

Scientific Prediction in Historical Sociology: Ibn Khaldun meets Al Saud¹

Peter Turchin
University of Connecticut

FINAL DRAFT: December 11, 2003²

One of the hallmarks of a mature discipline is its ability to make predictions that can be used to test scientific theories. Scientific predictions do not necessarily have to be concerned with future events; they can be made about what occurred in the past. I illustrate such retrospective prediction with a case study of conversion to Christianity in the Roman Empire. The bulk of the paper deals with the logic and methodology of setting up a scientific prediction in macrosociology. The specific case study I develop is the possible state collapse in Saudi Arabia. The theoretical setting is provided by the demographic-structural theory of state collapse. The starting point is a previously developed model for political cycles in agrarian societies with nomadic elites, loosely based on the ideas of Ibn Khaldun. I modify the model to fit the characteristics of the modern Saudi Arabian state and estimate its parameters using data from published sources. The model predicts that the sovereign debt of Saudi Arabia will reach unmanageable proportions some 5- 20 years in the future; the fiscal collapse will be followed by a state collapse in short order. The timing of the collapse is affected by exogenous events (primarily, fluctuations in world oil prices) and by parameter uncertainty (certain parameters of the model can be estimated only very approximately). The generalized prediction of eventual Saudi collapse together with subsidiary relationships specifying how variations in exogenous factors and parameters affect the future trajectory is the “Ibn Khaldun scenario.” A major theoretical alternative is provided by a set of ideas and specific recommendations suggesting how Saudi Arabia can avoid crisis by reforming its economy and liberalizing its political system (the “IMF scenario”). The main purpose of the proposed test, therefore, is to determine which of the two theoretical scenarios will best describe the trajectory of the Saudi state over the next two decades.

INTRODUCTION

Scientific prediction

The ability of a discipline to make valid predictions is a sign of its maturity (Collins 1995:1588). As historical sociology matures, its practitioners have begun to wrestle with the question of prediction. A good illustration of this preoccupation is the Symposium on Prediction in the Social Sciences at the 1993 meeting of the American Sociological Association (see Hechter 1995). The participants and commentators of the symposium expressed a wide range of opinions on the possibility of predicting such macrosociological phenomena as revolutions, ranging from

¹ I would like to thank Bertrand Roehner, Andrey Korotayev, Darya Khaltourina, Artem Malkov, Sohail Inayatullah, Tom Hall, and Edgar Kiser for helpful comments on the previous versions of this paper. Direct correspondence should be directed to Peter Turchin, Department of Ecology and Evolutionary Ecology, University of Connecticut, 75 N. Eagleville Rd, Storrs, CT 06269-3042, USA. E-mail: turchin@uconn.edu

² Submitted to the *American Journal of Sociology*.

“it cannot be done” to “it already has been done” (Kiser 1995). This disagreement is due, in large part, to differing definitions used by the authors. Thus, the first order of business is to clarify the various meanings of the word “prediction”.

The usual meaning of “prediction” is a statement that a certain kind of event will occur at some future time. What distinguishes prediction in science from the common usage is that we must have an explicit scientific theory on which the prediction is based. This requirement leaves beyond the pale “predictions” propounded by pundits at TV talk shows (no explicit theory) or astrological predictions (the underlying “theory” is unscientific). Within the scientific usage, we can further distinguish three kinds of predictions. The first, and conceptually the simplest one, is *projection*. In a projection exercise we ask a “what if” question: assuming certain initial conditions and a certain mechanism of change, what would be the future trajectory of the modeled system? An example is demographic projections that we can run for different scenarios of future fertility changes in the US. Whether the total fertility rate stays constant, declines, or increases will have a strong effect on the future age structure of the US population (e.g., Lee and Tuljapurkar 2001).

A *forecast* is a prediction that a certain variable will reach a specified level (or will be within the specified range of values) at a certain point in the future. Unlike the projection exercise, forecasting requires that we accept the validity of the assumptions of the underlying theory. A common example of a forecast is the weatherman on the TV predicting that the temperature will be between 70 and 75 degrees F at noon two days hence. Forecasts are made for a variety of practical reasons usually having nothing to do with science.

The third kind of prediction (I will call it *scientific prediction* to distinguish from the others) is used to test scientific theories. Scientific prediction inverts the logic of forecasting: whereas in making forecasts we assume the validity of the underlying theory and want to know what will happen to observables, in a scientific prediction exercise we want to use the observables to infer the validity of the theory. We take it for granted that theories yielding predictions that are in good agreement with empirical patterns are preferable to those who make poor predictions. The distinction I make here between forecasts and scientific predictions roughly parallels the distinction between unconditional historical prophecies and (also) scientific predictions, made by Karl Popper and endorsed by Michael Hechter (1995:1522).

I would argue that it is scientific predictions, rather than forecasts (and especially projections), that are the hallmark of a mature science. Why should macrosociologists care about forecasts? No practical person (e.g., a politician, or an intelligence analyst) should pay any attention to forecasts made by sociologists about, for example, the possibility of revolution in any particular state. At present time historical sociologists are only beginning to understand the causes of revolution and state collapse. Certainly, we are not at the stage where meaningful forecasts are possible. In other words, making forecasts is a useless activity given the current state of development of historical sociology. On the other hand, making scientific predictions could be a very fruitful activity, because the end result may be precisely what we currently lack – understanding. At least, that’s how it worked in natural sciences, and there is no reason to assume that it will be impossible to repeat this feat in social sciences (Turchin 2003).³

³ A lively discussion on the methods of explanation in historical sociology has developed during the 1990s (Kiser and Hechter 1991, Somers 1998, Kiser and Hechter 1998, Steinmetz 1998). In this paper I do not need to choose sides in this largely philosophical and highly abstract debate. My intent, rather, is to construct a specific and concrete case-study in scientific prediction, leaving the philosophical issues to philosophers.

Another strike against forecasts is that many mature sciences lack the ability to make accurate forecasts. Taking again the weather prediction, it is well known that no meaningful weather forecasts can be made farther in the future than 7- 10 days. The reason is that the dynamical system governing fluctuations in temperature, pressure, and rainfall (which we call “weather”) is characterized by sensitive dependence on initial conditions (popularly known as chaos). The characteristic time of trajectory divergence is such that very similar initial conditions, well within the observation error of the global network of weather stations, can and do produce completely different outcomes after just 7- 10 days. Although we cannot make long-term weather forecasts, the science underlying weather fluctuations (a discipline of physics known as fluid dynamics) is completely understood. Thus we have a seemingly paradoxical situation where the scientific understanding of processes is highly mature while our ability to predict long-term dynamics is nonexistent.

How should we use scientific prediction in macrosociology? Actually, setting it up is not a trivial task (which is why the ability to make scientific predictions is a sign of maturity in a discipline). The main requirement, of course, is a well-developed and sufficiently formalized theory from which to derive testable predictions. Ideally, the theory should be formulated in mathematical language, which makes it hard to sneak in hidden assumptions. Even more importantly, mathematical formalization increases the potential for quantitative predictions, which have a greater falsifiability potential. This consideration leads us to one thorny problem in scientific prediction, how to judge whether the theoretical statement to be tested has a high falsifiability potential, that is, whether it is *novel* (Elman and Elman 2003). Should I, for example, use any theory to make a prediction that the US will not have a state collapse during the next month, and then observe that, indeed, this is what happened, the reaction will be deservedly underwhelming. A successful, but trivial prediction should not have any confirmation value. But how do we judge which predictions are trivial and which are novel? The solution is simple: we should never test a theory in isolation, but always against alternatives. The simplest alternative could be a “null hypothesis,” an expectation that future events occur entirely at random. A better alternative is some sort of formalization of “common sense.” However, ideally we would like to be in a position to test our theory against a fully developed alternative theory, or theories (this is described more fully in Turchin 2003:7-8). Having two (or more) explicit theories allows us to determine for what aspects of reality their predictions coincide and, more importantly, where they disagree. Then, we collect data on the aspect where the theories disagree. The confirmation status of the theory, whose prediction matches the data better, is elevated, while the unsuccessful theory’s status is degraded (and eventually, after it makes a lot of poor predictions, we might decide to call it “rejected”). Thus, a well-designed exercise in scientific prediction can be defined as a planned comparison between two (or more) scientific theories using the data as arbiter.⁴ We now see that scientific prediction is a technically and conceptually much more difficult exercise than “mere” forecasting. Essentially, we must obtain multiple forecasts, from each of the theories we are testing. Furthermore, it may be necessary to run multiple projections to test the effects of assumptions (this will be illustrated below for the Al Saud prediction).

⁴ It is important to note that a prediction exercise is not the only way that general propositions (theories) can be tested. Other ways include checks for logical coherence, verification of the postulated proximal mechanisms and processes of change, examinations of comparative evidence (Richerson and Boyd 2001), analytic narratives (Bates et al. 1998), and at least two kinds of experiment, manipulative and mensurative. A prediction exercise is essentially a mensurative experiment, and only one way of many to do good science.

To sum up, in contrast to forecasts (as defined here), which are useless to the scientific progress, each successful exercise in scientific prediction advances our understanding of the processes involved. It is the ability to make scientific predictions that is a hallmark of maturity in a scientific discipline. It means that the discipline has accumulated enough theoretical models to erect explicit alternative hypotheses, and that it has developed structures for translating the chaos of raw sensory inputs into *data*, formalized empirical observations that can be used in testing theories (“normalized data” and “relevant facts” in the formulation of Rozov 2000).

Retrospective prediction: an example

One important aspect of the definition of scientific prediction given above is that there is nothing in it about predicting the *future*. Logically it does not matter whether the empirical event (or events) that we use to distinguish among theoretical alternatives will take place in the future, or has already occurred. The importance of this aspect cannot be overemphasized, because if we had to wait for our scientific predictions to be fulfilled (or not), the progress in historical sociology would be very slow, or even impossible (how could we test sociological theories about agrarian empires, if there are no such polities left in the world?).

As an illustration of such retrospective prediction, or retrodiction (Kiser and Hechter 1991), consider the case of testing dynamical theories about religious conversion (Turchin 2003). Three more-or-less explicit models for religious conversion and ethnic assimilation have been proposed in the literature: the noninteractive, the autocatalytic, and the threshold models. The justification for each of the model does not concern us here (the details are in Turchin 2003: Chapter 6); what is important is that each model predicts a qualitatively different trajectory (the proportion converted/assimilated as a function of time). This means that we can determine which theory better reflects the reality if we can find data on the temporal course of conversion. Empirical data on conversion to Islam in Iran and Spain, early Christianity in the Roman Empire, and the growth of Mormonism all strongly supported the autocatalytic model and were nothing like trajectories predicted by the two alternatives. What do we conclude from this result? All models are by definition wrong, because they oversimplify the complex reality, but the autocatalytic model is less wrong than the alternatives. We can tentatively conclude that the assumptions of the conversion process built into the autocatalytic model are approximately correct (at least, until an even better alternative model is proposed).

A critic might point out that the shape of the empirical conversion curves was already known to the tester prior to the test, and therefore the result of the test was a foregone conclusion. There is some validity in this criticism, so let us delve a bit into the issues involved. A good paradigm is provided by the distinction made in statistics between fitting models to data and using fitted models to *predict out-of-sample data* (“in-sample” refers to data used in model fitting, “out-of-sample” data are those that were not used in fitting but were reserved for testing the model; or perhaps were collected after the model was fitted). Indeed, flexible statistical models, given a sufficient number of parameters, can fit almost any imaginable shape. Thus, a true measure of how well a statistical model does at capturing certain aspects of reality can be obtained only by forcing it to predict out-of-sample data.

But the three conversion models that I considered were not flexible statistical models. They were based on specific assumptions about mechanisms underlying conversion, and predicted qualitatively different shapes of trajectories. Thus, the comparison between theoretically predicted shapes and the empirically observed ones was definitely a step forward, because it roundly rejected two of the models in favor of one. Nevertheless, the putative critic is

partially correct because successfully predicting out-of-sample data should always carry more weight than predicting data already in hand.

Actually, there was an element of out-of-sample prediction in the test involving the early Christianity data. This case study came from the book by Rodney Stark (1996) on the rise of Christianity (see also Hopkins 1998; Stark 1998). Stark used a variant of the autocatalytic model to predict how the number of Christians in the Roman Empire grew from the first century on. He estimated (guessed, really) that there were roughly a thousand converts in 40 CE and that their numbers grew at the rate of 40% per decade. Several years after he made these estimates, a colleague attracted his attention to the reconstruction by Roger Bagnall of the growth of Christianity in Egypt, based on data in Egyptian papyri. Since Stark was unaware of Bagnall's data at the time when he constructed his prediction, we have here a true test with out-of-sample data.

The story gets even better. Two years after I wrote the chapter on conversion in my book on Historical Dynamics I happened on a reference to a German dissertation that gave a list of Pagan and Christian office-holders between 324 and 455 (von Haehling 1978). I immediately realized that this data provide me with an opportunity to make another test of the theory. We can treat the Bagnall data as the "in-sample", on which the model parameters were fitted (and published prior to the knowledge of the von Haehling data). The von Haehling numbers, thus, are the "out-of-sample" data, on which the model's predictions are tested. The results are shown in Figure 1, where the predicted curve is calculated using the formula on p. 107 in Turchin (2003) and the parameter values as previously published: the initial proportion of the Roman population converted to Christianity at 40 CE, $g = 0.0017\%$ (p. 111) and the relative growth rate, $r = 0.034 \text{ yr}^{-1}$ (Table 6.1).⁵ We see that the curve fitted to the Bagnall data (showing the proportions converted before 300 CE) does a very good job predicting the course of Christianization in the von Haehling data (after 330 CE). The coefficient of prediction (the proportion of variance of out-of-sample data predicted by the model) is a very healthy 0.57. This is a remarkable result, given that the data are quite crude, affected by fairly large observation errors. We note that the predicted curve slightly overpredicts the data (5 data points above the curve compared to 10 points below). This is as should be expected – after all, the curve was not *fitted* to the data.

Why predict future?

I think that the above example illustrates very nicely that successful scientific retrodiction is possible and (more importantly) fruitful in historical sociology. Of course, a very suspicious critic might point out that there is a possibility of dishonesty on my part. Indeed, it is conceivable that I had the prior knowledge of the von Haehling data, and then craftily did not reveal it until my book came out, so I could claim to have made a successful out-of-sample prediction two years later. This is not what happened, but how could I prove it? This potential (if unfair) criticism, I think, reveals why we really value predictions about future as the strongest kind of test of our ideas and theories – it precludes cheating. In my opinion, we tend to overestimate the value of predictions about the future, and we do it not for logical, but for psychological reasons. After all, retrodiction can also be set up to preclude any possibility of cheating. For example, one could use the theory to predict some data that have not yet been unearthed by archaeologists.

My argument above should not be taken as an exhortation to completely avoid predictions about future events (just that we should not limit ourselves to them). Under certain

⁵ Note that r here stands not for the coefficient of correlation, but it is a parameter in the model.

circumstances, such predictions can be set up to yield valuable tests of theories. One particular example that can be used in a prediction exercise is Saudi Arabia, to which I now turn. I begin by discussing the theory to be tested, which can ultimately be traced to Ibn Khaldun. Next, I propose an explicit model, based on the general theory but taking into account the peculiarities of the Saudi setup. Third, I consider empirical sources for the estimation of parameters. Finally, I return to the general issue of scientific prediction (as opposed to forecasting), and discuss how the Saudi Arabian exercise helps us to clarify the meaning and limitations of the approach.

AN EXPERIMENT OF SCIENTIFIC PREDICTION: STATE COLLAPSE IN SAUDI ARABIA?

Theory: Ibn Khaldun

Abd-ar-Rahman Abu Zaid ibn Muhammad ibn Muhammad ibn Khaldun (1332- 1406) was a statesman, jurist, and historian, and perhaps the first sociologist in the modern sense (Gellner 1981). He is best known for his remarkable theory of political cycles, based on his intimate knowledge of Islamic societies of the Maghreb (Northern Africa west of Egypt). Ibn Khaldun's theory can be extended (naturally, with many modifications) to apply to agrarian societies in general (this is discussed in Turchin 2003: Sections 3.2.1 and 7.1). The ideas of Ibn Khaldun, furthermore, can be synthesized with theories of state collapse proposed by modern sociologists (the main influence is the work of Jack Goldstone, see Goldstone 1991) to develop a general theory of state collapse in agrarian societies (Turchin 2003).

The core of the theory, as it is currently formulated (Turchin 2003: Chapter 7, following Goldstone I call it the demographic-structural theory), concerns the relationship between population growth and fiscal stability of the state. Briefly, population growth in excess of the productivity gains from the land leads to persistent inflation and rising real costs, which outstrip the ability of the state to increase tax revenues. Rapid expansion of population also results in an increased number of aspirants for elite positions, putting further fiscal strains on the state, and intensifying intra-elite competition and factionalism. Increased rural misery, urban migration, and falling real wages lead to frequent food riots and wage protests; expansion of youth cohorts contributes to the population mobilization potential; and elite competition and popular discontent fuel ideological conflicts. As all these trends intensify, the end result is state bankruptcy and consequent loss of military control; elite movements of regional and national rebellion; and a combination of elite-mobilized and popular uprisings that manifest the breakdown of central authority (Goldstone 1991:25).

This verbally formulated theory has been formalized by constructing a suite of mathematical models covering various combinations of assumptions about the structure of studied societies. Of main relevance to the topic of this paper is the model that was appropriately named "the Ibn Khaldun's model" (Turchin 2003: Section 7.2.3). Here I give an abbreviated description of the model, while in the next section we will see how it can be adapted to the case of modern Saudi Arabia.

The dynamics of the Ibn Khaldunian "world-system" are determined by the interaction between the civilized society and the desert tribes. The civilized region is the site of recurrent state building/collapse episodes. It is inhabited by indigenous commoner population, who provide the productive basis of the society. The desert is inhabited by stateless (but not chiefdomless) tribes, who periodically conquer the civilized region and establish a ruling dynasty there. Desert tribes, thus, supply the ruling elites for civilized states. Initially the ruling dynasty establishes

government that is moderate in expenditures and just in administration. General prosperity results in population growth, and causes both rulers and people to become accustomed to increased spending (“luxury”). The army and bureaucracy demand and receive higher pay. As habits of luxury increase, and must be paid for, the state attempts to increase its revenue through heavy taxation, or outright seizure of its subjects’ property. This fiscal policy inevitably leads to the ruin of economy, followed by famines, pestilence, political unrest, and eventually state collapse. The area then is reconquered by desert tribes, who establish a new dynasty and the cycle repeats.⁶

In the model I simplify the Ibn Khaldun scenario by focusing on just two components of the system: the state fiscal health and the elite dynamics (note that such drastic simplification is a key step in building successful models; I will discuss this issue in more detail in the next section). I assume that the dynamics of commoner population are largely disconnected from the elite dynamics. Dynasties come and go, but peasants and merchants continue to grow food, trade, and pay taxes to whichever government is currently in power. (This assumption, of course, greatly oversimplifies the reality, and thus I investigated an alternative formulation that models commoner dynamics explicitly, see Turchin 2003:134). Thus, the rate of resource extraction from commoners is a constant, R . During the early years of the dynasty, the extracted resources are divided in two parts: taxes to support the government, and rents to support the elites.

Elite dynamics are characterized by two variables: their numbers and average per capita income. One important parameter in the model is μ_{\min} , the per capita income that nobles consider to be the minimum that accords with their station. This “minimal acceptable income” is determined socially, and can vary between societies. Ibn Khaldun argued that with time former tribesmen forget the rude ways of the desert, and subsequent generations grow accustomed to ever increasing luxury. Thus, μ_{\min} is a variable that starts at some low level at the beginning of the dynasty and then increases at a certain rate. Elite numbers increase as long as their incomes exceed the minimum. I estimated the maximum (per capita) rate of increase, $r_{\max} = 0.08 \text{ yr}^{-1}$ at four times the intrinsic rate of population increase typical for preindustrial populations.⁷

The key assumption of the model is that as long as income per elite capita, generated from rents, exceeds the minimum acceptable income, the state and elites live in harmony. However, if elite numbers grow to the point where their per capita incomes fall below μ_{\min} , then nobles become dissatisfied, and will use a variety of techniques to divert some of the taxes into their pockets.

The state fiscal dynamics are modeled as the balance of revenues and expenditures. Revenues consist of a fixed proportion of R as long as elites are not too numerous. Eventually the numerical growth of nobility leads to the decline of state revenues. The expenditures are proportional to the elite numbers, because elites demand employment from the state as army officers, bureaucrats, and courtiers. Thus, the dynamics of S , the accumulated state resources, follows a typical trajectory through time. During the early period of the dynasty S grows, because elite numbers are few and their appetites are modest. At some point, however, the revenues drop to the point where they cannot match expenditures, and S declines, and eventually become 0. At this point, the model assumes that the dynasty failed: it is abandoned by the army and civil officials whom it can no longer pay. The state becomes vulnerable to conquest, which (at least in

⁶ A very important part of Ibn Khaldun’s theory is how the *asabiya* (group solidarity) of the dynasty waxes and wanes during the cycle, but here I do not have space that discussion of this fascinating topic requires (see Inayatullah 1997, Turchin 2003: Chapter 3).

⁷ This estimate is based on the assumption that the legal maximum of wives that a Muslim man could have is four.

the model) happens immediately, because the desert tribes provide the ready and spatially adjacent source of the next dynasty. A typical trajectory predicted by the model is illustrated in Figure 2. Numerical investigation of the effect of parameter values on the dynamics of the Ibn Khaldun model indicated that the main parameters that affect the period of the cycle are the maximum rate of elite population increase (r_{\max}) and the rate at which the minimal acceptable income grows with time. Rather rapid cycles of about one century in period, shown in Figure 2, obtain for high values of r_{\max} that should be typical for societies where elite polygyny is widespread.

The test case: Al Saud

Modern Saudi Arabia, of course, differs in many important respects from the medieval Maghrebin societies, which provided the inspiration for Ibn Khaldun's theory and the model described above. Surprisingly, however, we can tailor the model, described in the previous section, to the situation of Saudi Arabia without needing to add much complexity. Furthermore, as we shall see, the basic dynamics predicted by the model carry over to the modern case.

First, Saudi Arabia is not an agrarian society. However, it has a greatly simplified economy which allows us to easily modify the theoretical model to fit the peculiar conditions obtaining there. Saudi economics is dominated by the oil sector, which currently accounts for about a third of its GDP (SAMA 2002). The role of oil in the government revenues is even greater – currently about 80% (SAMA 2002) – and much of the rest of revenue is indirectly tied to oil. As a result, fluctuations in oil prices provide a very reasonable predictor of government revenues (see below). In short, we can assume a single-sector economy without much loss of accuracy. Thus, R of the original model, interpreted as the product of commoner labor, becomes the oil-derived state revenues. An important modification of the Al Saud model is to make R a variable that fluctuates from year to year in response to changes in oil prices.

Second, in the place of a single class of desert-originating elites we need to put something that is a better approximation of the complex structure of the modern Saudi society. Ideally, we would keep track of the numbers, rates of increase, and consumption levels of at least four different classes:

1. The royal family, consisting of the descendants of the founder of the Saudi dynasty – this is the House of Saud (or Al Saud, hence the name of the model).
2. A small group of commoner businessmen closely associated with the dynasty (example: Osama Bin Laden's father). These individuals have higher incomes and reproductive rates than some of the royal princes.
3. Government employees whose livelihood depends directly on the state budget.
4. The rest of the commoner population, whose salaries derive from other sources than the state budget, but who nevertheless depend on it for subsidized health services, education, and consumption items (such as food and fuel).

Unfortunately, the empirical sources available to me do not allow parameter estimation for all these classes. Accordingly, I compromised by lumping the first two into “elites” and the last two into “commoners” (note that since all four classes have a claim on a share of R , they together correspond to the “elites” of the Ibn Khaldun model). Furthermore, since the data about the wealthy associates of Al Saud are even harder to come by than for the royal family (see below) I will further approximate this stratum with just the royal family.

A quick digression on the purposes of modeling

Given this background, we now can write the equations of the model that we will use to predict the future dynamics of the Saudi Arabian polity. Before I do this, however, it would be a good idea to discuss the general logic of the modeling approach I use, because the use of formal models is just beginning to penetrate social sciences. The most important thing to keep in mind when constructing a model is that one should not aim to capture the reality in all of its glorious complexity. Experience from many fields of science shows, over and over again, that putting too much complexity in models defeats their purpose and leads to scientific failure of disciplines that insist on doing so. A good model should include only those processes that are critical to making predictions about the output variable (or variables). Of course, when we embark on a modeling exercise, we do not yet know which processes are critical. To find out, we begin by writing the simplest possible model and then sequentially add various candidate processes to it, at each step testing whether the addition of a process has a substantial effect on the predicted trajectory. Any process that has only a slight effect on model predictions, is ruthlessly expunged from the model, even if we perfectly well know that it operates in the real world. To repeat, the purpose of the predictive model is not a faithful description of the reality, but identification of the key processes one really needs to make accurate predictions about the system trajectory.⁸ Therefore, I ask my readers to temporarily suspend their disbelief as they read through the description of the model in the following paragraphs. The equations provide a starting point of the investigation, and the influence of other, non-modeled, processes will be discussed in due time.⁹

The equations of the Al Saud model

The “output variable” that we are focusing on is S , the state accumulated surplus (or deficit; in fact, the key question is how the state debt will grow with time). To write the actual equations for S I use a discrete-time formulation because much of the data come as yearly numbers (this approach will also make it easier for others to reproduce my calculations, should they wish to do so; to facilitate such checking I am posting an Excel file with calculations on the Web).¹⁰ The guts of the model are the equation governing the accumulated state surplus (deficit) in year t , S_t :

$$S_{t+1} = S_t + R_t - \mu_c C_t - \mu_e E_t \quad (1)$$

where S_{t+1} is the next year’s value of S (the quantity we are predicting), R_t is the annual state revenue (a function of oil prices), C_t and E_t are the numbers of commoners and elites, respectively, and μ_c and μ_e are the annual average income from the state that commoners and elites expect. I will assume that the numbers of commoners and elites grow exponentially:

$$C_{t+1} = r_c C_t$$

where r_c is the per capita growth rate of the commoner population. The equation for the elite numbers is the same, but the growth parameter is r_e .

⁸ This true in the context of scientific prediction; in other contexts different kinds of models, including purely descriptive, may be appropriate.

⁹ Western (2001) provides a very useful perspective on using the Bayesian framework for finding the balance between overly simple and overly complex explanations.

¹⁰ Available on the web at <http://xxxx>.

Parameter estimation

The next step is to obtain estimates of parameters. My main source of quantitative data is the annual reports by the Saudi Arabian Monetary Agency (e.g., SAMA 2002). There are some obvious problems in relying on the official data, but as far as I could determine, the numbers are reasonably honest and check against independent sources. The main problem with the official source, as we shall see later, is that it omits some critical information (especially that relating to the finances of the royal family). In Figure 3 I plot some of the critical variables that can be used to estimate model parameters: the Saudi state's revenues and expenditures, population numbers of Saudi Arabia, and the annual fluctuations of oil prices. Saudi finances throughout this article will be given in terms Saudi Riyals (USD = SR 3.75).

We need to make sense of these raw data. The first observation that we make is that there is a clear connection between the price of oil and total state revenues. The relationship is quite strong, with a linear regression using oil price as the independent variable explaining over 80% of variance in the state revenues (Figure 4a). Note that the dependent variable in this regression is the total annual revenue (derived from both oil and other sources; in fact, "other" revenues also fluctuate in response to oil prices). Thus, we have a simple equation predicting Saudi revenues,

$$R_t = cO_t$$

where O_t is the price of a barrel of oil in dollars, R_t is the Saudi revenue in billions of Riyals, and c is the regression coefficient (the slope of the straight line in Figure 4a). The regression estimate of $c = 8.43$ SR billion · barrel/USD.¹¹

State expenditures exhibit a somewhat more complex pattern of fluctuations, because first they are affected by oil prices (via revenues, when the state is flush with revenues the tendency is to spend more, when times are tough, some belt-tightening occurs). But there is also a second important influence that is revealed when we plot *per capita* expenditures (total expenditures divided by population numbers). Per capita expenditures experienced an enormous growth during the heady days of the 1970s (Figure 4b), when the total revenues of the state skyrocketed and the social contract between the House of Saud and the Saudi population was forged. When the oil prices collapsed in the mid-1980s, the per capita expenditures declined only to a threshold of around SR 10,000. Thereafter, good years (in terms of oil prices) saw upward movement of per capita expenditures, but during the bad years, Saudi rulers ran into difficulties whenever they tried to decrease per capita expenditures below the threshold of about SR 10,000 (more on this below). If this argument is correct, then we have an estimate of the parameter μ_c , the average annual income that commoners have become accustomed to receiving from the state. Thus, $\mu_c \sim$ SR 10,000. (Remember, that this is not a direct handout, but rather a combination of salaries paid to officials and various subsidies that make life easier for everybody.)

The final parameter that the SAMA data yields is the per capita rate of commoner population growth, r_c . Currently it is between 3 and 3.5% per year. Thus, we have reasonably solid estimates of the model parameters relating to the commoner module. What about the elites?

The Saudis are very closemouthed about the internal arrangements within the ruling family. Even the size of the family is a closely kept secret. As a result, I have encountered wildly different estimates in the published literature (mostly news articles). The most reasonable

¹¹ The weird-looking units of c are a result of the need to balance the units of R_t and O_t .

guesstimates appear in a web publication by one of the firms specializing in geopolitical and geoeconomic analysis (STRATFOR 2000).¹² STRATFOR analysts estimate the size of the royal family at between 10,000 and 20,000 members. Reportedly, the minimum allowance, received by the lowest ranked princes is \$50,000 per year. The stipends of higher-level princes are in the neighborhood of \$1- 2 mln, but for those at the top just monthly expenses can run into many millions of dollars. Perhaps the most relevant number is the STRATFOR estimate that about 10% of the Saudi budget goes to the allowances for the ruling family. Another estimate is that the House of Saud's annual budget is around 15% of the national income (Aburish 1995:294). Unfortunately, this expenditure does not appear in the SAMA numbers. According to Saïd Aburish (1995:295), the greater part of the royal budget is taken out of the oil income before it is recorded as national income.

I will assume that there are currently 10,000 princes and that on average they receive \$1 mln (SR 3.75 mln) per year. Thus, the total expenditure on the royal family is SR 37.5 bln, which is about 13% of Saudi revenues – right in the middle of the two available estimates.

It is similarly difficult to estimate the rate of population increase characterizing the House of Saud. Some numbers that crop up repeatedly in various publications (which, however, should not add to their credibility – there is a well-known tendency by news reporters to repeat “factoids” over and over again) is that when there were 5,000 princes the family grew at 35- 40 princes per month, implying the annual growth rate of 8- 10%. This level seems reasonable in light of the estimate of 8% per year for generalized Islamic elites proposed in the previous section. Thus, let $r_e = 0.08 \text{ yr}^{-1}$.

A “quick-and-dirty” test of the theory

Before we use the model to make projections, we need to run some quick-and-dirty checks on whether Saudi trajectory up to date is consistent with the theoretical predictions (actually, retrodictions). The first variable to check is the state's accumulated surplus, S . As we saw in the previous section, the Ibn Khaldun model predicts that S should increase during the early phase of a new dynasty, reach a maximum, and then decrease. The Al Saud model predicts a similar qualitative pattern (the quantitative details, such as the timing and the magnitude of the peak in S , will depend on how population growth and consumption level parameters change with time). Since we have the time series of annual revenues and expenditures, it is a simple matter to calculate how the state's accumulated surplus/debt evolved with time (this is the variable S_t). As Figure 5a shows, the qualitative shape of the observed trajectory is precisely as predicted by the model. The estimated accumulated surplus reached a peak of SR325 bln in 1982, and declined to negative SR425 bln by the year 2001. Independent information is in general agreement with this estimate. It is reported that during the early 1980s, the Saudi surplus was close to \$100 bln (SR 375 bln). The reported public debt in 2002 is SR 650 bln, somewhat greater than the estimated (SR 425 bln). Probably the difference is due to the accumulated interest on the public debt (I could not find any references to interest payments in the SAMA publications; this reticence is probably due to the religious prohibition of usury in Islam).

But could this boom-and-bust pattern be due simply to fluctuating oil prices, rather than effects of population growth? To disentangle the effects of oil prices from that of population

¹² Incidentally, the STRATFOR publication reaches conclusions very similar to those presented here, although their argument is wholly qualitative, not being backed up by an explicit model.

growth on S we need to run some explicit projections. I employed a simplified version of Equation (1), which omitted the elites term:

$$S_{t+1} = S_t + R_t - \mu_c C_t$$

Starting the projection in 1970 with $S_{1970} = 0$, I assumed that $R_t = 8.43 O_t$ (as estimated above, using the historical oil price data for the O_t series). Then, I approximated $\mu_c = \text{SR } 2,500$ for the period prior to 1975 and $\text{SR } 15,000$ for the period after 1975. Finally, for C_t I used two scenarios: exponential growth with the historical average of $r_c = 0.04 \text{ yr}^{-1}$ and no growth ($r_c = 0 \text{ yr}^{-1}$). Comparing the two trajectories (Figure 5b) we see that the critical factor is definitely population growth. The projection of the state surplus assuming population growth at the observed rate shows the boom-bust cycle (the quantitative details of the timing and the height of the peak in S are not captured accurately, but that is not surprising, since we used a simplified model). By contrast, if there is no population growth, then the state surplus continues to rise (oil prices affect the steepness of the rise, but not the qualitative pattern of dynamics).¹³

In other words, the Saudi state trajectory is generally on track postulated by the demographic-structural theory. What other social variables can we check? One of the most important mechanisms leading to the state collapse, as identified by Goldstone, is the rise of aspirant elites and increased inraelite competition. Goldstone proposed that we can measure inraelite competition by the number of individuals seeking higher education. The number of students enrolled in universities and colleges has been growing steadily (Figure 6). Even before the early nineties the rate of growth was quite impressive, but after c.1994 it literally exploded. This “credentialing” revolution (Collins 1979) suggests that recently the competition for higher-level jobs has intensified. Indeed, the major employer of college graduates is the state. Between 1965 and 1985, the number of civil servants increased tenfold (!), which allowed the state to absorb almost all university graduates. After 1985, however, the bureaucratic build-up reached the saturation point, with the result that educated Saudis had to take jobs below their level of competence (Aburish 1995:99). Currently, the unemployment rate among college graduates has been increasing. In general, there is clearly a severe case of “elite overproduction”. The new Saudi class of rich merchants, bureaucrats, teachers, doctors, and officers in the armed forces has increased from 2% to 11% of the population by the late eighties (Aburish 1995:100). This numbers, of course, refer to the middle-rank elites; the growth of the high ranks (essentially, the House of Saud) has been commented on earlier.

Other indicators of enhanced sociopolitical stress include the crime statistics: the number of reported crimes increased from 1775 in 1966 to 21,826 in 1985 (Aburish 1995:100). A crime wave, such as that experienced by Saudi Arabia recently, appears to be one of the fairly reliable indicators of the coming demographic-structural crisis in historical data (Fischer 1996).

Projecting model trajectory forward in time

The discrepancy between the calculated and actual debt level in 2002 raises an important point: as the Saudi debt reaches serious proportions, we can no longer neglect the cost of servicing it. Thus, the actual model that I will use in projecting S_t is

$$S_{t+1} = S_t + R_t - \mu_c C_t - \mu_e E_t + gS_t \quad (2)$$

¹³ According to the model, if the Saudi population had increased at the rate of 2% per year (instead of 4%), then the government could have kept its budget balanced during the thirty year period.

where g is the interest rate (the plus sign in front of the gS_t term means that when S_t is negative, the interest is subtracted from it). The mid-1990s data on Saudi debt servicing in Table 5.5 of Wilson and Graham (1994), themselves based on IMF estimates, suggest that $g \sim 7\%$, a reasonable enough rate for a country that can guarantee its debt with oil revenues.

Using Equation (2) together with equations for the growth of commoner and royal population it is now a simple matter to project S_t forward in time, using the starting 2002 value $S_{2002} = \text{SR } 650 \text{ bln}$. The Excel file (Al Saud.xls) which accompanies this paper provides a handy tool for accomplishing this projection. Interested readers can also use this spreadsheet to investigate how changing different parameter values (see below) affects the outcome. If we assume parameter values as estimated above (these are our *median* parameter values) and the average price of oil equal to \$25 per barrel (since the stated goal of OPEC is to keep oil price in the corridor of \$22- 28), then we observe that the Saudi debt grows with time at an accelerating rate: the projected debt exceeds SR 1 trillion by 2005, and SR 2 trillions by 2010. Clearly, at some point the state will not be able to secure further financing to cover its mushrooming deficit. Let us say, for concreteness sake, that the credit dries up when the annual interest payment exceeds the annual state revenue.¹⁴ The projections suggests that bankruptcy, defined in this way, will take place in 2014.

So here we have a concrete number – 2014. Is this the prediction of when Saudi Arabia will experience state collapse? No, because 2014 is just a single number, and meaningless by itself. One of the reasons for exposing the inner workings of the model that lead to this number is to show how tenuous some of the parameter estimates and other assumptions are. Therefore, the next step in constructing a proper scientific prediction is to investigate the effect of various sources of uncertainty on the predicted trajectory. There are three general sources of uncertainty: stochastic exogenous effects, uncertainty associated with parameter estimates, and the effect of structural assumptions of the model. I discuss each of these sources in turn.

Exogenous stochasticity

No model can incorporate all possible factors affecting the real-life processes we study and attempt to predict. The factors left out of the model are called the exogenous effects, and we usually model them as stochastic variables. In our model stochastic variables include fluctuations in population growth rate (e.g., due to an epidemic), interest rates varying with response to global economic conditions, and (most importantly) fluctuating oil prices. In fact, it is so obvious that oil prices will have a dominant effect on the projected trajectory, that I will focus on them exclusively (there is also a possibility of very large-effect – catastrophic – perturbations to the system, but I will deal with them when discussing structural assumptions). In fact, we have already done most of the work, by estimating the relationship between annual oil price and Saudi state revenue. Thus, it becomes a simple matter of investigating various scenarios concerning future oil prices. Here I will simply bracket the estimate by assuming constant very low or very high oil prices. First, it is conceivable that oil prices could fall as low as \$15 per barrel and stay there for a decade. If this happens, then the collapse date advances to 2008. By contrast, high oil prices of \$35 per barrel postpone the day of reckoning to 2020. We see now that the point estimate of 2014 gave us a quite false feeling of precision. Taking uncertainty in just one factor –

¹⁴ Actually, the financial collapse will probably occur prior to this, but I cannot think of any other simple stopping rule.

oil prices – yields an interval of 13 years; according to the model, state collapse could occur during any year between 2008 and 2020.

Parameter uncertainty

The next important source of uncertainty is the estimates of parameters. The first thing we should investigate is the effect of the ruling family, because there was so much uncertainty associated with the estimates of the family size and per capita allowances. We can do it in one fell swoop by first, reducing the effect to zero, and second, by doubling it. Surprisingly, the effect is not as strong as might be expected (Table 1). If we completely cut the princes off the feeding trough, we only delay the collapse by three years, to 2017. Doubling the number of princes (or, alternatively, doubling the estimate of average allowance per prince) advances the collapse date to 2012. There are two important messages here. First, the effect of uncertainty associated with the estimates of princely parameters is substantially less than that associated with future oil prices. This means that there is no particular need to spend a lot more effort trying to sharpen these estimates. Simply put, the main dynamic is driven by the growth of the commoner population.¹⁵ Second, and the corollary to the first, although the luxurious life style of many Saudi princes excites much (deserved) opprobrium among the commentators, and more importantly the Saudi population, our calculations show that they are only a small part of the problem.

The next set of parameters to investigate are those for the commoner population: their population growth rate and the average income they get from the state. In principle, both of these rates can decrease with time. Population growth rate in Saudi Arabia was recently close to 4%, and now it has apparently declined below 3.5%. Suppose that it will linearly decline by 1% each decade. Surprisingly, neither this assumption, nor a more extreme one of 2% decline per decade, shift the collapse date at all (Table 1). Taking the expectations of government subsidy next, suppose that the government succeeds in persuading the populace to accept a linearly declining μ_c at the rate of SR400 per year (that is, in ten years μ_c will decrease from SR10,000 to 6,000). This would serve to move the collapse date to 2018. By contrast, an increased demand on the state resources of the same magnitude, but opposite direction (resulting, for example, from the necessity to create government jobs to fight unemployment) would shift the date forward to 2012 (Table 1). Finally, the interest rate that the government has to pay to service the debt has a similar – moderate – effect on the trajectory (Table 1).

The effect of structural assumptions: Ibn Khaldun versus IMF

This category of uncertainty in the prediction is the most interesting one from the scientific point of view. In fact, finding out whether structural assumptions are correct is what the empirical test is all about. The issue of structural assumptions can be approached by putting it in the broad context, that is, by considering the theoretical alternatives to the demographic-structural model. We can start delineating alternatives with the New York Times column of Thomas Friedman on February 27, 2002, entitled “One Country, Two Futures”. In it, Friedman describes two possible models for Saudi Arabia’s future. The first one predicts collapse, using the reasoning broadly consistent with the model developed above. The second model proposes an alternative to

¹⁵ This argument should not be taken as invitation to sloppy parameter estimates. Rather, one should devote efforts towards greater precision where it makes a large difference. The formal approach that allows one to do this is called sensitivity analysis (Saltelli et al. 2000).

collapse by assuming that it is possible to reform the Saudi system both economically and politically. The recipe for modernization is utterly familiar to anybody who follows the news on international political economy. It involves opening up the economy to external competition (which among other things means entering the World Trade Organization), balancing the budget by cutting social spending, and liberalizing the political system with the ultimate goal of rule of law, democracy, and free elections. Since this is the standard package pushed upon the developing world by such organizations as the International Monetary Fund (IMF), let us call it the “IMF alternative”.¹⁶

Both the IMF and the Ibn Khaldun alternatives are based on the logic of *endogenous* development of the Saudi state. Yet Saudi Arabia does not exist in isolation. Its future is affected by the actions of other regional powers, including (in addition to itself) Egypt, Israel, and Iran (Iraq being now out of this equation), superpowers (of which there is currently only one), and non-state networks, such as Al Qaida. There is a possibility that one of these exogenous actors will exert an overwhelming force, which would in a sense “spoil the experiment”. One potential scenario, which has already been mooted in the US policy circles, is that the US might decide to intervene militarily in Saudi Arabia by taking over the oil-rich littoral. The rationale of such an intervention might be to overthrow the evil Saudi regime that promotes Islamic terrorism worldwide; or, alternatively, to prop up the friendly Saudi regime, our bulwark against the tide of Islamic revolution in the Middle East. The opposite possibility is the American failure in Iraq, leading to the rise of an aggressive Islamist (most likely Shiite) regime there, which would then export Islamic revolution to the Shiite areas of Saudi Arabia (who predominate in, again, the oil-rich littoral). One can multiply the scenarios, but these two or three should suffice. I do not know how probable any of them are, but the important point is that we do not need to have this information. This is because we are concerned here not with a forecast, but with setting up a scientific prediction. Any of the scenarios mentioned above, let us call them “the exogenous intervention” alternative, would result in “spoiling the experiment”, which concerns the prediction of how endogenous dynamics of Saudi Arabia will develop. Thus, what we have to do is to predicate the empirical test on the exogenous intervention *not happening*. If the international environment allows Saudi Arabia to develop endogenously, then we have our test, if not, we are out of luck – we will not know who is right in this particular case, Ibn Khaldun or IMF.

Returning to the two endogenous alternatives, we can use them to structure the discussion of structural assumptions of the Al Saud model.

Delivering significant economic growth. The most important assumption of the IMF scenario is that it is possible to stimulate substantial economic growth in Saudi Arabia. Since the population is growing very rapidly (at more than 3% per year), in order for per capita incomes to increase in

¹⁶ Here are the actual recommendations made by the IMF’s executive board in October 2002, as reported by the Middle East Economic Digest of November 1, 2002: implement a “comprehensive privatisation strategy”, “use part of the proceeds from privatisation to reduce the public debt”, “reduce the barriers to inflow of foreign direct investment”, approve and implement “the capital market and insurance laws”, “balance the budget by 2005”, and implement the income tax, with “sales tax as a good interim measure pending the implementation of a fully-fledged value-added tax (VAT)”. IMF recommendations are primarily addressing the economic aspects of reforms that Saudi Arabia is urged to pursue. For the political aspect, we can look to an Op-Ed piece in New York Times by Kenneth M. Pollack (October 16, 2003). Pollack is currently director of research at the Saban Center of Middle East Policy at the Brookings Institution and a former director for Persian Gulf affairs at the National Security Council. Commenting on a recent announcement by Saudi Arabia to hold elections for municipal councils with a year, Pollock wrote: “The only way for the Saudis to get at these deep-seated problems is through modernization, and that process has to start with the political system”.

a noticeable manner, the GNP needs to grow even faster. Furthermore, as a result of age structure dominated by youthful cohorts, the Saudi labor force is currently growing at 4.5% per year (Hatrash and Fareed 2002). In order for the country to go beyond the mere absorption of new workers, the long-term growth rate in GDP must be no less than 6% (Hatrash and Fareed 2002). The Ibn Khaldun alternative, at least as formalized in model (2), explicitly assumes 0% GNP growth in the non-oil sector, and no growth of oil-related GNP except that entirely driven by oil prices. In other words, the Saudi state will continue to rely exclusively on oil revenues. The truth is likely to be somewhere in between, and the empirical issue is whether substantial economic growth can be delivered, or whether it will be too slow for per capita incomes to rise appreciably (if not decline). The past record of IMF and the World Bank in delivering economic growth is not good (Easterly 2001; Naiman and Watkins 1999).

Cutting budget deficit. In the short run (as opposed to the possibility of long-term economic growth), this will have to be done by reducing government expenses and/or raising taxes. Since the bulk of expenditures goes to the salaries for government employees (60%) and social programs, any expenditure-reducing measures will immediately affect the quality of life for the Saudi population. It is doubtful that this can be done. The Saudi population has become accustomed to the state subsidies during the 1970s, when state revenues expanded as a result of upsurge in the oil prices (Figures 3 and 4b). When oil price collapsed in the early 1980s, the Saudi economy went into recession, and the state started running huge budget deficits. This forced the government to reassess the Kingdom's generous welfare program with its costly reliance on subsidies (Wilson and Graham 1994:185). The Saudi government was able to erase some subsidies on agricultural products. The government was lucky in that this action coincided with falling agricultural commodity prices (previous attempts to cut agricultural subsidies had run into stiff opposition). Other attempts at erasing subsidies (on gasoline, water, and electricity), however, were unsuccessful in the face of immediate protests. Plans to increase the price of gasoline and electricity, announced in 1984, were rescinded the next year. These setbacks forced the government to shift its track in 1987-8 and try to increase revenues. The revenue-increasing measures included new user fees (e.g., an airport departure tax) and imposing an income tax on repatriates. "Four days after it has been published, the expatriate tax was withdrawn. Two weeks later, most of the ... user fees were scrapped in the face of mounting domestic pressure" (Wilson and Graham 1994:186-9). Another unsuccessful attempt to reduce subsidies occurred in 1992. Again, the resulting outcry forced King Fahd to rescind his decision before it was applied and instead he had to increase some subsidies to placate the people (Aburish 1995:306). Currently (November 2003) despite the IMF calls for implementing the income tax, the Saudi government has repeatedly stated that it has no plans to do so (except on the expatriates, but so far it has not taken any concrete steps to do it). Thus, it is an empirical observation that the Saudi government has been unable to reduce the budget deficit, by any means, whenever oil prices collapsed. Each time these measures provoked significant public resistance and unrest, and had to be rescinded. Furthermore, it can be argued that the very foundation of the social contract on which the Saudi regime is based is the commitment to spread the oil wealth. Any movement away from continued social spending immediately undermines the legitimacy of the regime.

Reducing population growth. The present regime, given its dependence on conservative Muslim (Wahhabite) clerics, cannot propose any population-control measures. On the other hand, it is possible that population growth rates will decline on their own, as a result of deep social forces acting on individuals. It is a fallacy to assume that Muslim populations are incapable of collectively reducing their fertility rate – the best counterexample is the recent fertility collapse to the replacement level in Iran (and under a theocratic regime, to boot). On the other hand, the Iranian fertility decline occurred after more than a generation of heightened sociopolitical instability. It takes time for social mores to change. Furthermore, even if the Saudi fertility were

to be reduced to the replacement level tomorrow, it would still take multiple years before the actual population growth would cease. This is known as the demographic inertia effect, and again it takes about a generation for things to equilibrate. What all this means is that there is a substantial lag time (a generation or more) before any social changes will percolate to a reduced rate of population growth. According to the model calculations, Saudi Arabia does not have this time.

Liberalizing the political system. The Ibn Khaldun scenario assumes that the Saudi elites will continue to cling to power to the bitter end, and go down with the state, when everything collapses. This is of course what invariably happened in pre-modern times. The modern era, by contrast, knows a number of examples of peaceful power transition from autocratic to democratic regimes. The IMF scenario assumes that this is a realistic option for Saudi Arabia. Some kind of opening of the political system may in fact be necessary in order to push through unpopular, but needed reforms, the argument goes.

Actually, I would argue that the logic is inverse: autocratic (but perceived as legitimate) regimes seem to have a better track record in pushing through unpopular reforms (I admit, though, that I am not familiar with any quantitative studies on this subject). In any case, it is not clear at all that the Saudi elites would be willing to allow any liberalisation. After all, reforming the system is an extremely risky business for the elites involved, as the Shah of Iran and Mikhail Gorbachev found for themselves. The Saudi ruling class is quite aware of this danger.¹⁷

Furthermore, by implementing any changes in the political system, the Saudi elites will be cutting themselves off from power and wealth. This does not seem like a very rational strategy. A clear-headed Saudi prince might be completely aware of the collapse to come, in which case the most rational strategy is to continue milking the Saudi oil revenues for as long as possible, salt away most of the wealth abroad, and when the system collapse leave the country to live off the accumulated wealth. Any rocking of the boat may only hasten the inevitable end.

Buying time by selling assets. Opening up economy and privatizing government property, as some liberal economists urge, may result in one positive side-effect: using the proceeds from privatisation to finance the deficit. Such a strategy could, in principle, stave off the collapse by some years. In fact, we can even obtain some quantitative estimates. The state controls one-third of shares on the Saudi stock exchange, whose total capitalization is \$80 bln. Thus, the state share is SR 100 bln, or less than one-sixth of the 2002 deficit. In other words, selling the state controlled shares will only stave off the collapse by 2 months.

What about selling ARAMCO, the state-owned oil company? This is after all the main asset of the House of Saud. It is highly dubious that this is politically possible, but let us estimate (to the order of magnitude) how much money could be raised in this way. Saudi Arabia can produce over 10 mln barrels of oil per day (currently the actual production is under 10 mln b/d to keep the price up). At \$25 per barrel this is equivalent to just under \$100 bln. Thus, ARAMCO is worth anywhere between \$1 trillion (assuming a 10% return rate on investment) and \$2 trillion (assuming 5% return). This is a major amount of money, and if such a sale could be done, it

¹⁷ As reported by Elaine Sciolino in the New York Times of November 4, 2001. Prince Bandar bin Sultan, the Saudi ambassador in Washington told her about a letter exchange in the late 1960s between the Shah of Iran and King Faisal. The shah advised the king to modernize his country, because otherwise he could not guarantee that Faisal would stay on his throne. Faisal wrote back, thanking the shah for the advice, but refused to follow it. "History proved our point", concluded Prince Bandar.

would postpone the collapse by decades. Nevertheless, my guess is that selling off ARAMCO, whether wholesale or piecemeal, is a course of action that is politically unacceptable.

What are the consequences of fiscal collapse? One important assumption of the “Ibn Khaldun scenario” is that financial collapse of the Saudi government inevitably leads to the state collapse, followed by a combination of inraelite conflict and elite-mobilized and popular uprisings; in short, a protracted period of civil war. This is the typical trajectory in historical case-studies examined by Jack Goldstone (1991) and Turchin and Nefedov (2004), but it would be foolish to claim that precisely the same is going to happen in today’s world. After all, some modern democratic societies have experienced bankruptcy (the most recent example being Argentina), which resulted in the fall of the government but hardly civil war. There are reasons to doubt that the fall of the Saudi regime will lead to a peaceful and rapid transition to some other regime. In my opinion, just about the only remotely possible manner in which bloodshed could be avoided is a rapid takeover of the state by Bin Laden types. This is obviously a highly unpalatable outcome from the point of view of the Western world, and it is precisely this course of events that might trigger the intervention by the US. But this is just a guess, the main point I want to emphasize here is that the formal prediction of the Saudi trajectory stops at the point of fiscal meltdown. What happens after that point is beyond the scope of the exercise. Incidentally, this point of high trajectory indeterminacy following state collapse is discussed by Goldstone (1991).

Structural assumptions: a summary. One useful end result of considering the structural assumptions above is that we can see that the Ibn Khaldun and the IMF alternatives constitute the ends of a continuous spectrum. Thus, (and assuming that the exogenous intervention does not occur), it is conceivable that some sort of mixed scenario is what will actually happen. For example, there could be some moderate economic growth, Crown Prince Abdullah might succeed in reforming the system, the public may agree to some decrease in the state subsidies, etc. Will that be enough to prevent the fiscal followed by political collapse? The advantage of proposing an explicit model, as I did in this paper, is that all these questions are answerable. In fact, I have already investigated such effects as declining population growth or willingness to accept lower state expenditures on the part of population. Adding some terms to the model representing the effects of growing GNP is quite straightforward.

DISCUSSION

My main goal in this paper was to discuss the role that scientific prediction (which I distinguish from a mere forecast) can play in advancing the state of art in macrosociology. The bulk of the paper was devoted to constructing a case study in prediction, focusing on the possibility of state collapse in Saudi Arabia during the next two decades. I think most social scientists will agree that predicting a revolution is not a trivial task, in fact some feel it is outright impossible (Kuran 1995). Indeed, a credible *forecast* of Saudi Arabia’s future is impossible, given the current state of social science (if at all – see the discussion above about the impossibility of long-term forecasts in chaotic dynamical systems). On the other hand, I argue that we can set up an empirical test that will allow us to distinguish between predictions of two rival theories (the comparison that I termed “Ibn Khaldun versus IMF”). The strength of this approach, which I term scientific prediction, is that no matter what happens - the actual trajectory matching more closely the Ibn Khaldun or the IMF, or even some other unanticipated scenario - , we will learn

something useful (with the proviso that Saudi dynamics are allowed by external forces to develop largely endogenously).¹⁸

One of the key elements in making a scientific prediction is to publish the explicit algorithm on which it is based, so that others can use it to check on the internal logic and the empirical basis (e.g., parameter estimates) of the mechanism for generating predictions. Having an explicit parameterized model will be particularly important during the *post-mortem* stage of the analysis (once Saudi Arabia went into collapse, or alternatively after more than two decades have passed without it going into collapse). We can take it for granted that the future will bring some surprises and therefore the actual trajectory will be different from the predicted one. The job of the post-mortem analysis is to determine the source of deviations between the predictions and the reality. Some deviations do not challenge the theory. For example, it is impossible to predict how oil prices will behave – if I could do that, I would not be teaching undergraduates for living. During the post-mortem stage, however, we can simply substitute the observed time series of oil prices. Same with parameter uncertainty – many things (e.g., the number of royal princes) will probably come to light at some time in future, and better parameters can be substituted in the model in place of the preliminary estimates. But the real challenge (and the real test) is whether the model has correctly identified the main structural variables underlying the Saudi sociopolitical dynamics, and made the correct structural assumptions about how these variables affect each other. The main goal of the post-mortem analysis is to identify which of the structural assumptions held and which did not.

It is clear from the preceding paragraph that I consider the post-mortem analysis a key part of the whole exercise. Lack of a *quantitative* post-mortem analysis is perhaps my main critique of one of the most successful cases of prediction in historical sociology, Randall Collins' prediction of the Soviet collapse. The geopolitical model of Collins is eminently quantifiable (see, for example the explicit models by Hanneman, Collins and Mordt 1995; Turchin 2003). It would be a laborious, but doable exercise to estimate the geopolitical burden carried by the Soviet Union, and the additional strain of the Afghan war, which is when the tipping point was reached in the Collins scenario (Collins 1995:1568). And, incidentally, it would be interesting to see whether the US is in a similar position, given its current involvement in Afghanistan, and most importantly Iraq. Does the geopolitical theory predict the US decline? The data needed to answer this question are widely available.

The point of my criticism is not to belittle the achievement by Randall Collins; in fact, his prediction of the Soviet collapse stands out as the best example to date. The alternative theories, such as that of an impervious totalitarian colossus, have been decisively put out of the picture. But in order to have further progress, we need to shift from qualitative theories to ones that can be translated into explicit quantitative models. For example, I am not convinced yet that the geopolitical angle is the whole, or even the most important part of the story of the Soviet collapse.¹⁹ We need to test it against the, for example, demographic-structural alternative (and others; now that the Soviet Union has collapsed, the number of explanations for this event are proliferating). It is imperative to expand the number of worked-out case studies, because a mature theory cannot stand on a single empirical case, no matter how wonderful. The number of case

¹⁸ Personally, in fact, I hope that the Ibn Khaldun prediction fails, since I wish nothing but good to the people of Saudi Arabia, and hope that they will be able somehow to avoid the misery and suffering of civil war.

¹⁹ Note, also, that I did not use the geopolitical theory of Collins as an alternative in the Al Saud case study, because doing so would be creating a patent "strawman". There are no indications, as far as I know, that present-day Saudi Arabia suffers from an imperial overstretch in any form.

studies need not be huge – perhaps half-a-dozen or ten would do: Iran to supplement Saudi Arabia, Yugoslavia to supplement the Soviet Union, one-two from Latin America, Southern Europe, and Africa. I believe that once we have such a sample, general insights would emerge almost inevitably, and we will have a mature sociological theory of state collapse.

References:

- Aburish, Said K. 1995. *The rise, corruption and coming fall of the House of Saud*. New York: St. Martin's Press.
- Bates, Robert H., Avner Greif, Barry R. Weingast, and Jean-Laurent Rosenthal. 1998. *Analytic Narratives*. Princeton: Princeton University Press.
- Collins, R. 1995. "Prediction in macrosociology: the case of the Soviet collapse." *American Journal of Sociology* 100:1552-1593.
- Collins, Randall. 1979. *The credential society: an historical sociology of education and stratification*. New York: Academic Press.
- Easterly, William. 2001. *The elusive quest for growth: economists' adventures and misadventures in the tropics*. Cambridge, MA: MIT Press.
- Elman, Colin, and Miriam Fendius Elman. 2003. "Lessons from Lakatos." Pp. 21-68 in *Progress in international relations theory: Appraising the field*, edited by Colin Elman and Miriam Fendius Elman. Cambridge, MA: MIT Press.
- Fischer, David H. 1996. *The great wave: price revolutions and the rhythm of history*. New York: Oxford University Press.
- Gellner, E. 1981. *Muslim Society*. Cambridge: Cambridge University Press.
- Goldstone, J A. 1991. *Revolution and rebellion in the Early Modern world*. Berkeley, CA: University of California Press.
- Hanneman, R A, R Collins, and G Mordt. 1995. "Discovering theory dynamics by computer simulation: experiments on state legitimacy and imperialistic capitalism." *Sociological Methodology* 25:1-46.
- Hatrash, Hasan, and Saleh Fareed. 2002. "Kingdom faces three economic challenges." *Saudi Gazette*: January 21, 2002.
- Hechter, Michael. 1995. "Symposium on prediction in the social sciences. Introduction: reflections on historical prophecy in the social sciences." *American Journal of Sociology* 100:1520-1527.
- Hopkins, Keith. 1998. "Christian number and its implications." *Journal of Early Christian Studies* 6:185-226.
- Inayatullah, Sohail. 1997. "Ibn Khaldun: The strengthening and weakening of asabiya." Pp. 25-32 in *Macrohistory and macrohistorians*, edited by Johan Galtung and Sohail Inayatullah. Westport, CT: Praeger.
- Kiser, Edgar. 1995. "What can sociological theories predict? Comment on Collins, Kuran, and Tilly." *American Journal of Sociology* 100:1611-1615.
- Kiser, Edgar, and Michael Hechter. 1991. "The role of general theory is comparative-historical sociology." *American Journal of Sociology* 97:1-30.
- . 1998. "The debate on historical sociology: rational choice theory and its critics." *American Journal of Sociology* 104:785-816.
- Kuran, Timur. 1995. "The inevitability of future revolutionary surprises." *American Journal of Sociology* 100:1528-1551.
- Lee, Ronald, and Shripad Tuljapurkar. 2001. "Population forecasting for fiscal planning: issues and innovations." Pp. 7-57 in *Demographic change and fiscal policy*, edited by Alan J Auerbach and Ronald D Lee. Cambridge: Cambridge University Press.
- Naiman, Robert, and Neil Watkins. 1999. "A survey of the impacts of IMF structural adjustment in Africa: Growth, social spending, and debt relief." Washington, DC. URL: <http://www.cepr.net/IMFinAfrica.htm>; Center for Economic and Policy Research.
- Richerson, Peter J, and Robert Boyd. 2001. "The evolution of subjective commitment to groups: a tribal instincts hypothesis." Pp. 186-220 in *Evolution and the capacity for commitment*, edited by R M Nesse. New York: Russel Sage Foundation.

- Rozov, N S. 2000. *Approaches to a Rational Philosophy of History: The Research Program, Models, and Hypotheses. Volume 5 (In Russian)*. Novosibirsk, Russia: Novosibirsk State University.
- Saltelli, Andrea, K. Chan, and E. Marian Scott. 2000. *Sensitivity analysis*. Chichester: Wiley.
- SAMA. 2002. *Saudi Arabian Monetary Agency Thirty-Eight Annual Report*. Riyadh, Saudi Arabia: Saudi Arabian Monetary Agency, Research and Statistics Department.
- Somers, Margaret R. 1998. "'We're no angels': realism, rational choice, and rationality in social sciences." *American Journal of Sociology* 104:722-784.
- Stark, R. 1996. *The Rise of Christianity: A Sociologist Reconsiders History*. Princeton, NJ: Princeton University Press.
- Stark, Rodney. 1998. "E contrario." *Journal of Early Christian Studies* 6:259-267.
- Steinmetz, George. 1998. "Critical realism and historical sociology: a review article." *Comparative Studies in Society and History* 40:170-186.
- STRATFOR. 2000. *Saudi's Arabia future stability questionable*: www.stratfor.com.
- Turchin, Peter. 2003. *Historical dynamics: why states rise and fall*. Princeton, NJ: Princeton University Press.
- Turchin, Peter, and Sergey Nefedov. 2004. *Secular cycles*: Book manuscript.
- von Haehling, Raban. 1978. *Die Religionszugehörigkeit der hohen Amtsträger des Römischen Reiches seit Constantins I. Alleinherrschaft bis zum Ende der Theodosianischen Dynastie*. Bonn: Ph.D. dissertation.
- Western, Bruce. 2001. "Bayesian thinking about macrosociology." *American Journal of Sociology* 107:353-378.
- Wilson, Peter W, and Douglas F Graham. 1994. *Saudi Arabia: the gathering storm*. Armonk, NY: M.E. Sharpe.

Table 1. The effect of various assumptions in the prediction when Saudi Arabia goes bankrupt. (Median prediction = 2014). "Early" refers to the value of parameter that advances the date of collapse; "late" the opposite value that delays collapse.

Parameter	Early	Late
Oil price per barrel (median \$25): early = \$15, late = \$35	2008	2020
The drain of prince allowances on the budget: late = 0, early = double the median	2012	2017
Population growth rate of commoners: early = declines by 1% per decade, late = by 2% per decade (median = no change)	2014	2014
Change in μ_c : early = increase of SR 400 per year, late = decrease of SR 400 per year (median = no change)	2012	2018
Interest rate: early = 10%, late = 5% (median = 7%)	2010	2018

Figure Captions

Figure 1. A test of out-of-sample prediction of the proportion of Christians in the Roman Empire.

Figure 2. Trajectory predicted by the Ibn Khaldun model.

Figure 3. Financial and population data from the Saudi Arabian Monetary Authority (SAMA).

Figure 4. Analysis of the SAMA data.

Figure 5. Financial and social indicators in Saudi Arabia. (a) Estimated surplus/debt. (b) Two projections of the state surplus/debt: with population growth ($r_c = 0.04$) and without.

Figure 6. Financial and social indicators in Saudi Arabia, continued: the number of students in the institutions of higher education.

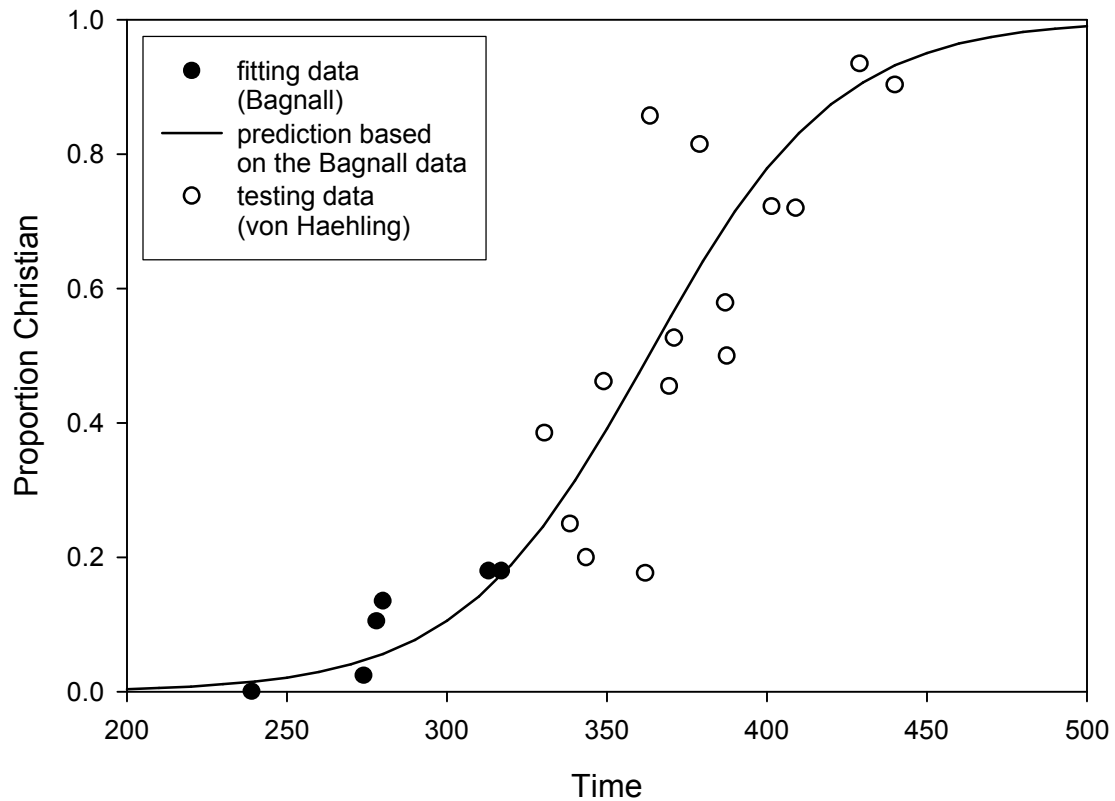


Figure 1.

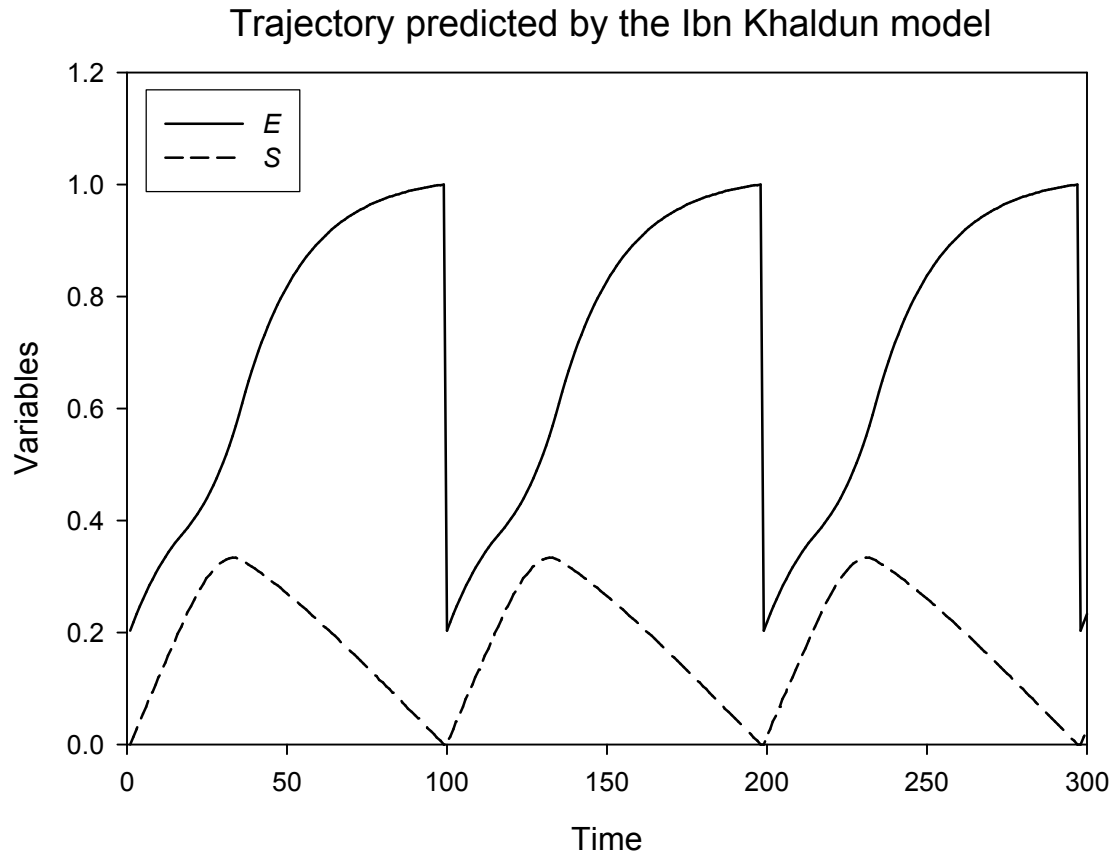


Figure 2

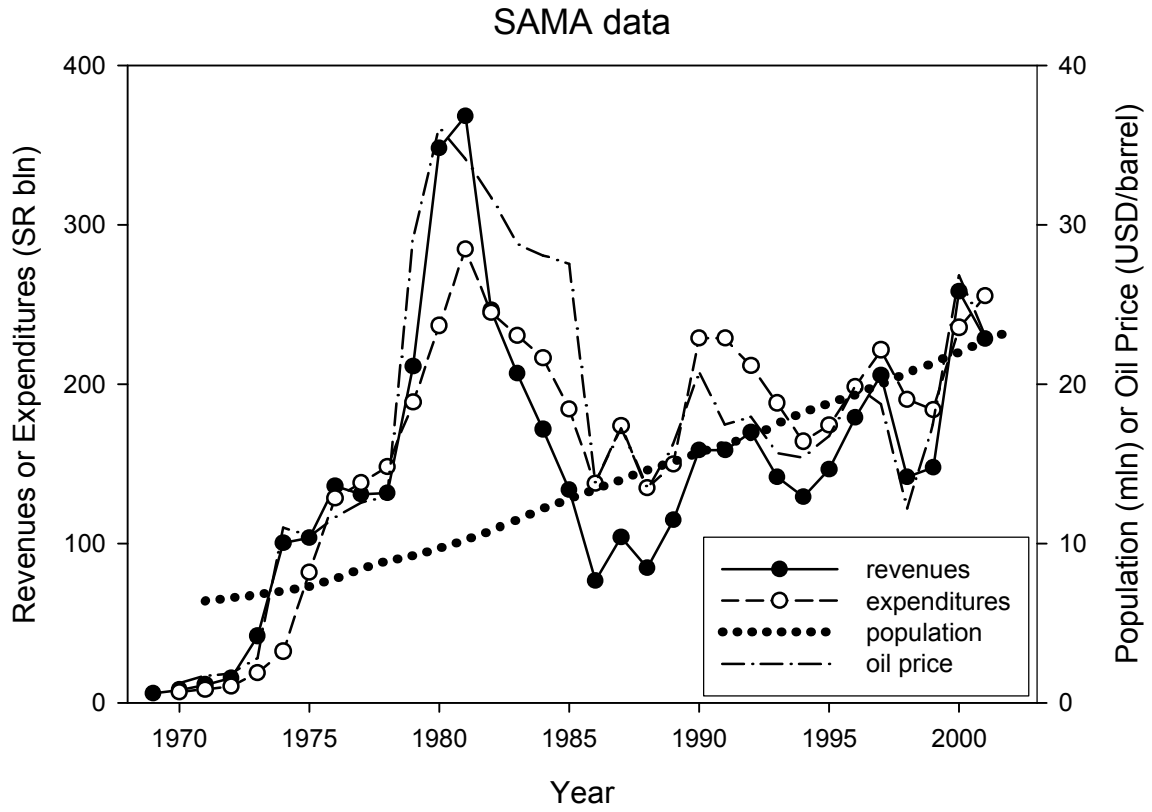


Figure 3

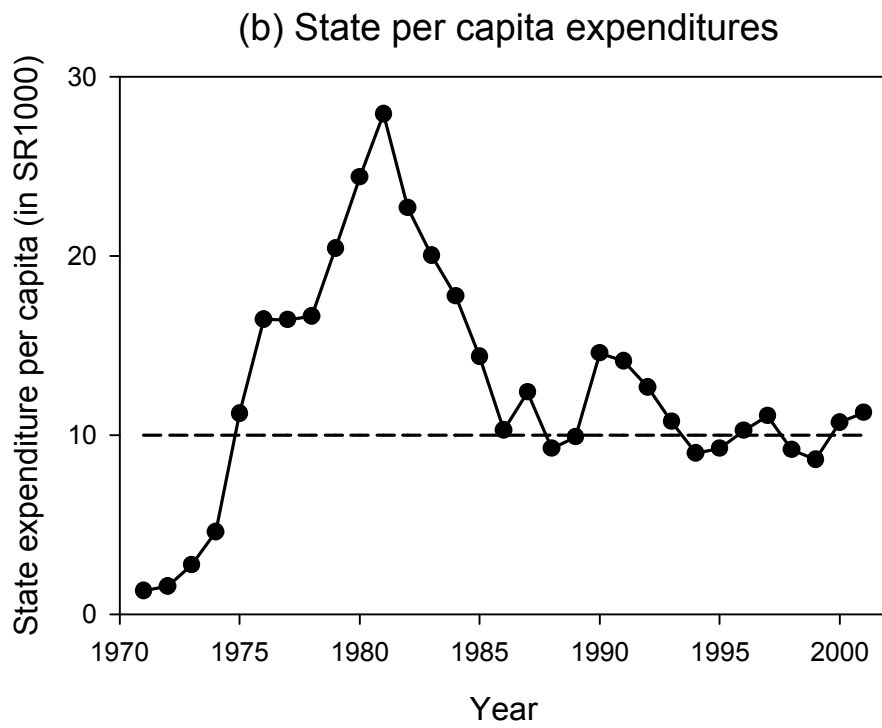
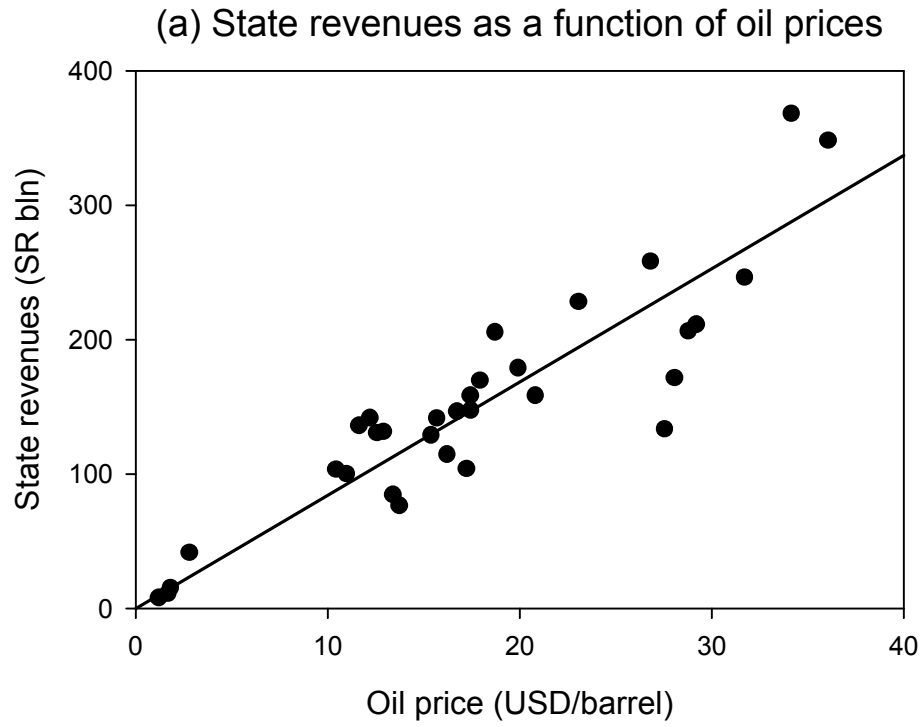


Figure 4

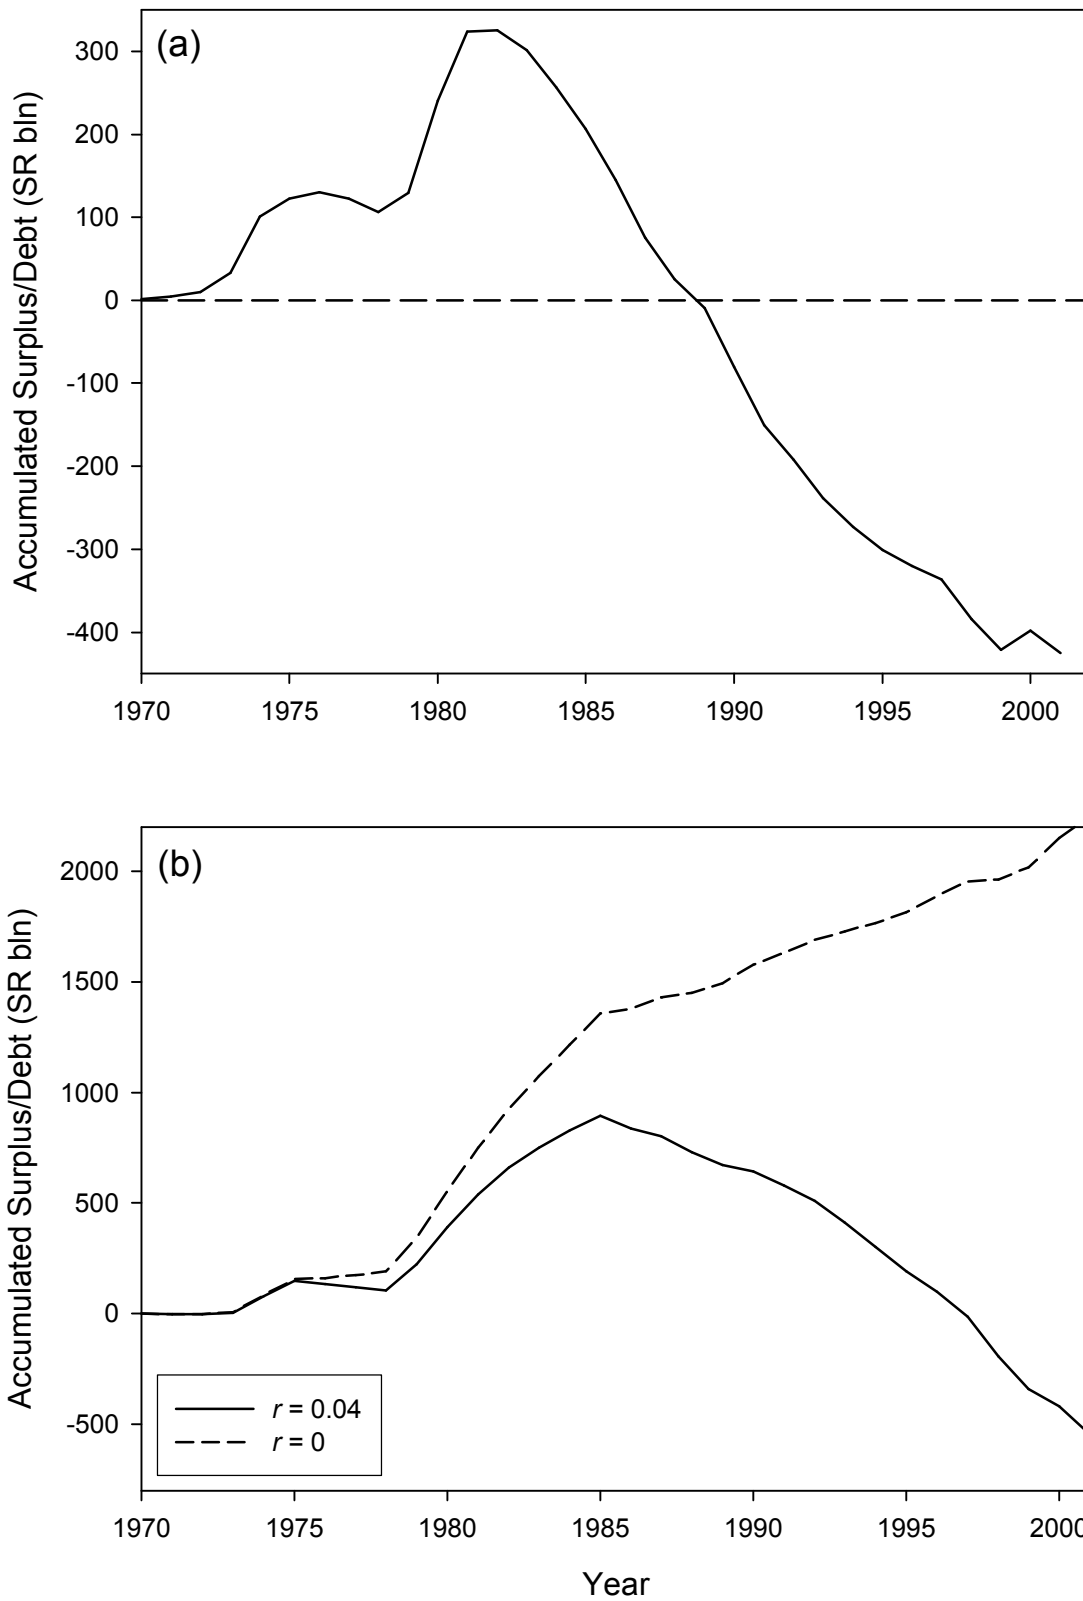


Figure 5

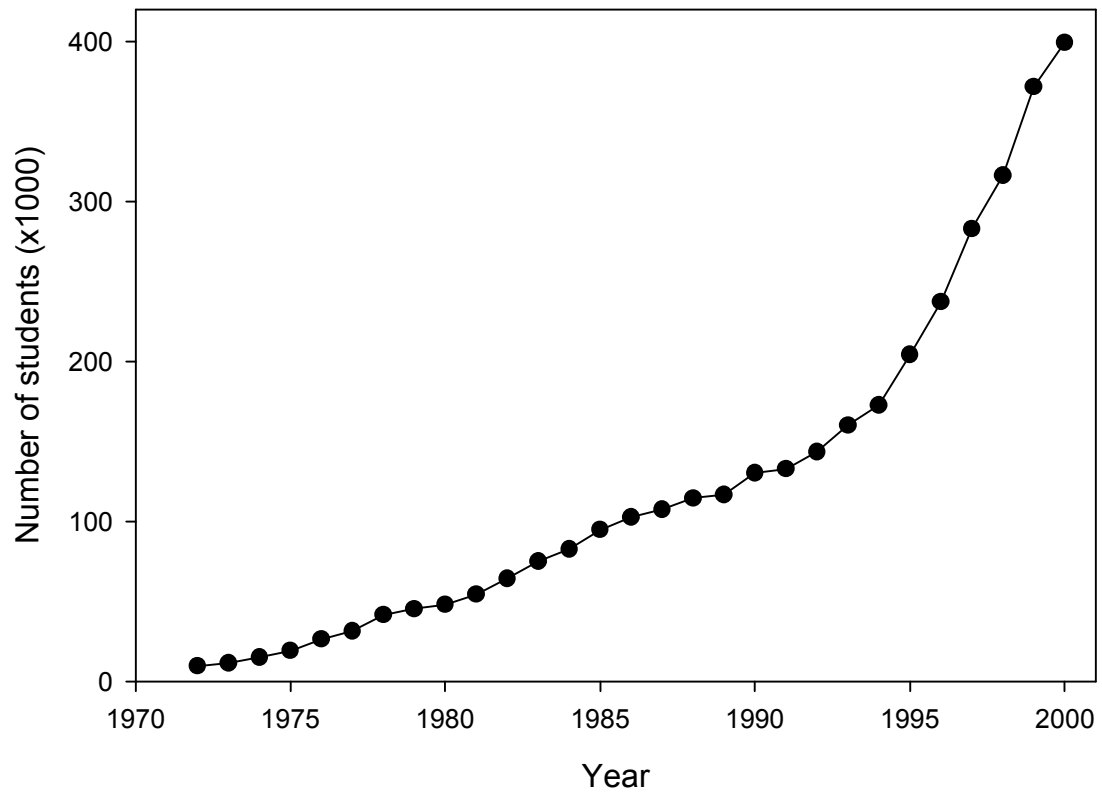


Figure 6