



Fall of Berlin Wall in 1989 signaled end of Cold War and decline of communism

DOD (F. Lee Cockran)

Reorienting Grand Strategy

The Promise of Single-equilibrium Defense Planning

By TIMOTHY J. JUNIO and JONATHAN PROTZ

Since the end of the Cold War, the U.S. defense community has focused its futures analysis on a “range of possible outcomes” approach. Planners assume that social behavior, such as that of states in the international system or individuals in markets, is so complex that it defies point prediction. The best one can hope for, goes this mindset, is for a creative mind to envision scenarios that might come to pass, and then to prepare capabilities and strategies to meet challenges in those notional worlds. This approach to planning neglects two key and undeniable facts. First, at a specified level of granularity, there will be only one outcome of social interactions under study—a single equilibrium—just as there will be only one state of reality 5 minutes, 5 hours, and 5 years from now. Understanding which equilibrium will result is an informational, not a logical, research problem. Second, the United States has enormous potential to affect and effect changes in its favor—that is, to drive social systems, such as the international system, toward the particular equilibrium that U.S. policymakers desire.

Timothy J. Junio is a Doctoral Candidate in Political Science at the University of Pennsylvania and an Adjunct Researcher with the RAND Corporation. He previously worked for the Office of the Secretary of Defense and the Intelligence Community. Dr. Jonathan Protz is Assistant Professor of Mechanical Engineering and Materials Science at Duke University. He has served as a Consultant to the Office of the Secretary of Defense and was an American Association for the Advancement of Science Defense Policy Fellow.

The traditional view remains correct—for now—regarding the lofty challenge of point prediction for most kinds of social systems. The persuasiveness of this perspective is eroding, however, due to radical improvements in the ability of the United States to acquire and analyze information and the potential for these improvements to make single-equilibrium strategy a preferred

the United States has enormous potential to drive social systems, such as the international system, toward the particular equilibrium that U.S. policymakers desire

approach. Thinking in terms of the range of possible outcomes distracts planners from efforts to achieve a particular outcome. By focusing on maintaining the stability of one ideal equilibrium—what one may call the lost art of grand strategy—defense planners would improve the probability that their desired equilibrium is attained. Such

exogenous shocks and possible interactions in the development of its strategy. Holding material capabilities equal, the deepest thinking actor's desired equilibrium is the one most likely to be attained.

In this article, we advocate regrounding U.S. defense planning in single-equilibrium terms. Given the current tenor of political debate in Washington, this approach to strategy may sound novel. In fact, it is a return to the grand strategy tradition of the United States during World War II. At that time, senior U.S. leaders chose policies to shape the Western political system such that the United States would emerge at its apex after the war. Applying this kind of planning process to today's challenges involves focusing on overarching objectives for the international system while continuing the process of alternative futures planning.

Our framework reconceptualizes scenario-based analysis as a means to return a social system to a desired path in the face of exogenous shocks. The United States could thus create a "funneling effect" on the future of the international system or subsystems of interest; bringing its enormous material

to single-equilibrium strategy is not only desirable, but also is a *necessity* given advances in technology. In particular, persistent surveillance, large-scale digital data retention, and advanced algorithms offer state actors enormous potential to better understand and, ominously, manipulate the behavior of social systems. We argue that applying these technologies to national security policy will become a competitive process between states during the coming decades, a claim that has an amoral descriptive component—the possibilities created by these new technologies—and strong normative implications regarding the changing relationship between states and their constituent populations.

One World or Many?

The national security community is plagued by a tension: should planners attempt to predict the most probable state of social systems of interest, or focus instead on the range of possible outcomes? Each approach has relative advantages and disadvantages. Point prediction of social systems allows national security policymakers to allocate their limited time to a particular contingency. Examples abound, but for the sake of illustration consider the potential for revolution in countries of interest to the United States. Political scientists and the Intelligence Community alike have found it extremely difficult to gauge with precision when and where social revolutions are likely to occur. However, attempting to quantify the probability of revolutions and thus rank-order states at risk of political disturbances has strong intuitive appeal for those on the National Security Council or in the Department of State, who would be responsible for dealing with regional crises. An alternative approach—planning scenarios—has an advantage of covering a large range of possible outcomes in the social system of interest, and thus has a high chance of offering analysis of the actual outcome that results. This could mean thinking in terms of what kinds of governments might come to power following a revolution in a country of interest—during the Cold War, this would clearly have been a planning mechanism for a state's fall to communism, while at present the United States may be concerned more about nationalist or theocratic regimes.

The difference in outlook afforded by each approach is significant. Those who prefer the scenarios-based approach argue that the historical record of social scientists



During World War II, President Roosevelt pursued grand strategy that would benefit postwar United States

Library of Congress (U.S. Signal Corps)

a reconceptualization of futures analysis focuses not on attempting to predict the future, but on treating alternative futures as the consequences of exogenous shocks to a single-equilibrium trajectory. The "deepest thinking" of strategic actors is the one that has accounted for the greatest number of such

capabilities to bear, U.S. shaping efforts may constrain the choices of adversaries and thus reduce the number of possible outcomes. Preparing for exogenous shocks, therefore, may occur in a narrower range and with the intent of returning a system to the preferred equilibrium. Additionally, we argue that the return

seeking to predict social behavior is ridden with false judgments. At the top of the list is the general failure of U.S. policymakers to predict the demise of the Soviet Union; those favoring scenario analysis would note that the Soviet Union's collapse was *anticipated* while not *predicted*. The process of anticipating consequential political outcomes and planning for them would, it is argued, better prepare policymakers for future decisions than would dedicating analytic resources to trying to pinpoint exact equilibria. Those favoring point prediction, on the other hand, note that delivering a list of future contingencies that are treated as probabilistically equal is not particularly useful to a policymaker working 18-hour days in the Pentagon or at Foggy Bottom. That person would never have the time or adequate knowledge to conceive of and plan for every possible contingency.

Those who focus on the inability of researchers—or governments—to effectively model social systems tend to note three kinds of limits: observational (data may only be collected on a small part of a system at a time because, otherwise, the sensory, computational, and retention requirements exceed the researcher's capacity); cognitive (even if sensors hypothetically could capture large parts of systemic interactions, researchers would be unable to understand the nature of the interactions); and psychological (modeling *social* behavior is constantly plagued by humans' annoying tendency to break with expectations of what the "correct" actions are per a utility function). The number of factors affecting the outcome in even the smallest social system is deemed so numerous, and makes possible such a large number of combinations, that overcoming these three obstacles to even *describe* a social system is pronounced a practical impossibility.

We offer the perspective that these analytic approaches may be combined in a grand strategy framework. Those focused on point prediction of social systems are correct to note that, in fact, there will be one outcome of current processes of social interaction. The extreme of this position was once stated by mathematician Pierre-Simon Laplace. He proposed a thought experiment of a hypothetical God-like entity who, if able to know the position and velocity of every particle in the universe, could predict the entire future. This thought experiment captures the difference between informational and logical hindrances to point prediction

in social systems: if sufficient data could be obtained and analyzed, a precise outcome could be determined at a reasonable level of granularity. The level of granularity refers to a model's level of abstraction in space and time; since space and time are infinitely divisible, researchers must determine the level at which to gauge whether or not a system of interest is in a stable equilibrium. For example, in terms of the international system, one may think of the distribution of power or states' interests as one level of granularity, and one year as the unit of time. In a smaller-scale social system, such as analyzing a particular state, the level of granularity would be much finer. One may, for instance, look at the preferences of clusters of individuals (for example, ethnic groups) in week-long periods. The long-term objective of social modeling on any system of interest would be to incrementally improve the model's granularity.

One might say this approach is about taking the assumptions out of economic analysis. If an information collection and retention system were sufficiently effective as to identify an individual's preferences, then why bother using a deductive approach such as utility functions? Those on the opposite end of the scenario planning spectrum, however, would rush to point out the aforementioned limitations of identifying even the current processes of social interaction, much less determining the result of those processes a year into the future. The relevant question dividing these approaches, then, is about what advances in data gathering and analysis are plausible regarding social system modeling.

Technology and Predicting Social Behavior

The emergence of several new technologies has placed mankind on the precipice of major breakthroughs in the ability to overcome observational, cognitive, and psychological limitations to predictive modeling of social systems. Critics consistently point out that despite the increasing availability of new information technologies, researchers have not gotten much closer to achieving social prediction. The core flaw in these types of criticisms is a failure of imagination. With an unlimited time frame, incremental improvements in technology will allow researchers to capture greater amounts of social behavior until a point is reached at which prediction is possible with a reasonable degree of confidence. Many decades may pass before social

system prediction is possible on a large scale, but we argue that actors with the long view in mind may gain first-mover advantages by beginning to develop technologies and policy processes to support single-equilibrium strategy in the present.

Consider, for instance, a monitoring system that observes, records, and analyzes the behavior of individuals in a pedestrian plaza sized a hundred yards square. Such a system could use cameras, hard drives, and an algorithm to study an aspect of the social system, such as the average amount of time a person spends sitting. The value of such a system could be, in its initial stages, understanding traffic flow for purposes of public safety. The first instantiation of the system could have one sensory input, one means of data retention, and one processing mechanism (the algorithm differentiating between the background environment and the individuals moving within it). Such a system, skeptics of social modeling might say, would do little more than a city employee sitting on a bench in that pedestrian plaza with a clipboard.

One may imagine, however, incremental improvements in the system over long periods of time that have an aggregate effect of creating a powerful surveillance tool. The algorithm may be modified to identify objects other than human beings—for instance, an unattended piece of luggage. Further software modifications could include identifying a specific individual of interest with gait or facial recognition. Another incremental improvement, in the memory capacity of the system,

the ability to automatically identify long-duration, repeated visits by the same person outside the White House or a U.S. Embassy would be of great interest

could allow for tracking the frequency of an individual's visits to the area. While this kind of monitoring system may be of limited usefulness at a dog park, the ability to automatically identify long-duration, repeated visits by the same person outside the White House or a U.S. Embassy in a high-threat country overseas would be of great interest.

The listed examples all refer to the *scope* of the system's monitoring capabilities. One

may liken this to increasing the number of variables in an equation; the system seeks to describe more of the social behavior within a given environment. All of the listed incremental improvements are well within the range of existing technology. An example of an improvement in the scope of the described system would be the ability to monitor individual preferences, such as for marketing purposes. For instance, an advanced algorithm could tag individuals in a commercial environment, such as a shopping mall, and develop models of consumer behavior by studying patterns of which stores individuals frequent. A further incremental improvement could be to remotely estimate the age of an individual (a quality related to gait and facial appearance) to add another level of complexity to the model. Yet another incremental improvement to such a monitoring system could include changes in sensory inputs; for instance, in a system designed to study consumer behavior, one would be interested in credit transactions. Incremental improvements in the availability of data, such as tracking sales statistics for those stores, would give the owner of the monitoring system a unique appreciation for the overall behavior of the shopping mall—which stores are most visited and by whom, and which have the highest rates of sales per shopper. More importantly, such a system would allow for the development of an inductive model of consumer preferences in a manner dramatically different from standard approaches such as survey data. Indeed, why study a survey drawn from a sample when the entire population of interest may be monitored?

The other means by which incremental improvements may lead to revolutionary breakthroughs in social system modeling is in the *scale* of monitoring systems. Our example was first limited to a small pedestrian area of a hundred yards square. A system may incrementally increase in scale, say, by another 50 yards every year or so. Material constraints (for example, how many cameras the local government could afford to buy and the number of analysts to parse the data) may dictate the extent to which the scale of the system may be extended when continuing to apply existing technology.

New technology also offers the ability to expand the scale of social monitoring, however. Consider an incremental improvement in the resolution of the cameras covering the area; such an improvement would allow for coverage

over a wider area as well as make possible qualitative improvements within the system (for instance, the new data create a new demand on algorithms able to analyze it). Such a system, using currently available commercial technology, could provide surveillance of the public social behavior of an entire city. Extremely high resolution imagery, such as gigapixel photography, allows an observer to take a picture of an entire city skyline and zoom in on an individual office window or restaurant. The use of this level of resolution, particularly if combined with aerial assets (such as blimps, unmanned aerial vehicles, or simply putting the camera at points of high elevation throughout the city), allows for an extreme extension of the scale of the system.

These examples help to stimulate one's imagination of what is possible, and the analytic distinction of increasing the *scope* and *scale* of a monitoring system helps to generalize our perspective regarding long-run incremental improvements. One more way to organize thinking about technological improvements to social system modeling is to focus on the kinds of technologies themselves, thus providing an analytic target to those seeking to understand what incremental improvements may finally lead researchers to attain confident prediction. The most relevant technologies, as noted in the above example, are persistent surveillance, digital memory storage, and advanced algorithms. Advances in each of these three types of technology are the key to understanding the potential for a revolution in the observation and prediction of social systems.

Persistent surveillance (PS) systems are those designed to maintain a constant, watchful eye over a target (for example, an individual person or a physical space, such as a large urban area). Among the most important technologies enabling PS is high-resolution imagery, such as gigapixel photography that allows a surveillant to zoom in on specific objects of interest from within a very large viewing frame.

Large-scale data retention refers to the wide availability of massive amounts of inexpensive digital memory. Google, for instance, draws on this type of information technology (IT) to archive the Internet. Combining large-scale data retention with PS technology allows surveillants to archive enormous amounts of data regarding a target of interest. In the above example of an individual under a multisensory PS system, large-scale data

retention would allow surveillants to develop models of that individual's behavior.

Algorithms are the "rules" that tell a computer what computations to perform and how to perform them. While software algorithms have necessarily existed since the beginning of the IT revolution, recent advances create new potentials to automate tasks normally performed by humans. In the case of PS systems, software performs many critical functions, including:

- generating models of target behavior, ranging from those of an individual to patterns within groups
- spotting outliers from an established model of "normal" behavior
- identifying targets of interest, such as by recognizing a person by gait or face, noting the use of "red flag" spoken or written words, or warning of the presence of dangerous objects based on shape or material composition
- discriminating between targets
- coordinating many sensors within an integrated surveillance system.

Taking these technologies together, one may see how incremental improvements during the next several decades will lead to unprecedented levels of understanding of social behavior. As described in our notional scenario, the extension of sensors in a surveillance network would allow for a high degree of understanding of how a particular individual behaves or, more ambitiously, to understand general tendencies within a large group of people. This level of knowledge, provided development of algorithms proceeds apace with sensory expansion and data retention, may finally allow for near-real-time

the extension of sensors in a surveillance network would allow for a high degree of understanding of how a particular individual behaves or, more ambitiously, understanding general tendencies within a large group

prediction of social behavior. An important assumption underlying this analysis is that while individuals may have knowledge of the surveillance systems under discussion,

for most persons, deviation from norms of behavior to avoid surveillance will be excessively costly. For instance, as evidenced by tepid responses to the explosion of closed-circuit television monitoring, far more than a majority of a population is likely to ignore surveillance systems in public places. While personal privacy concerns may abound, once surveillance systems are in place, most citizens have few incentives to attempt to escape observation by the system.

Implications

Once governments begin to achieve high levels of knowledge of social behaviors of interest, this knowledge may be used in combination with tools of national power to effect behavioral changes. Persistent surveillance systems are useful in helping governments understand what actions to take (learning the preferences of an individual or population of interest) and whom to take the actions against (knowing which and/or how many persons must be influenced to achieve a desired political effect). Traditionally, these tools of national influence have been military or economic. Increasingly, IT plays a role in the U.S. ability to shape desired outcomes. While the technologies described above portend revolutionary changes in the ability to *understand* foreign behavior, information networks also allow the United States to *influence* persons or populations around the globe. The most obvious examples of these communication technologies are various tools on the Internet (email, Voice over Internet Protocol) and cellular telephones. Governments may also influence foreign populations by punishing them via the content of information systems. The most prominent example is banking, which is almost completely digitized.

The development of precision weaponry is a useful analogue to nonviolent coercion of individuals over communication networks. During World War II, communications were indiscriminately directed at entire populations, such as with radio broadcasts and dropping leaflets from planes. Similarly, munitions were used against entire urban areas. In the following decades, however, “carpet bombing” gave way to the development of laser-guided, and later global positioning system (GPS)-guided, bombs. The increased precision of these weapons has resulted in the scaling down of their size; the military’s ideal is to avoid collateral damage by using the minimum amount of explosive.

Communication technology has followed a similar pattern whereby access to information networks allows for the precise targeting of an individual of interest from thousands of miles away. Whereas once a government may have focused on posting a Web page to influence an entire foreign population, now one may imagine a smaller “bomb” in the form of email directed to a country’s subpopulation of interest, or Web pages that display different content depending on the location of a user seeking to access it. More ambitiously, one may imagine in the near future a replacement of economic sanctions against an entire country—which may have unfortunate humanitarian side effects—with unilateral targeted freezes of individuals’ bank accounts.

One may consider these sorts of technological changes and their coercive potential from the perspective of game theory. There is no such thing as “perfect information,” although this assumption drives much of economic modeling. In the real world, the costs of acquiring and processing information, and cognitive biases and limitations, place some ceiling on the amount of data being incorporated into strategic decisions. The surveillance systems discussed here suggest a manner in which a strategic actor may begin to raise that ceiling.

Game theory allows for modeling decisions with imperfect or asymmetric information; this approach is fruitful for considering how advanced IT may affect the choices of strategists. Consider two strategic actors. Each seeks to understand the payoff structure of the other as precisely as possible. Discerning the rank-ordering of someone’s preferences is, in practice, often difficult—particularly since individuals generally do not sit around making decision trees and ranking their relative valuation of commodities or activities. By

observing an individual’s behavior, however, a researcher may begin to create a model “as if” that person’s preferences were known. The primary logic underlying this approach is similar to why we believe surveillance systems may work to achieve social prediction: deviating from preferred behavior for the sake of thwarting a surveillant is costly to an individual. One may think of this in terms of traditional explanations of collective action problems. A society may, in the aggregate, have an interest in deceiving a surveillant, such as its own government (under a dictatorship) or a foreign government (under a competitive vision of global information networks). For one person, however, the perceived payoff of rebellious behavior is imperceptible, while the cost—even if quite small—will be greater in nearly every circumstance.

The strategic actor with superior information may use that advantage to shape the decisionmaking context of the other actor, thereby driving the outcome of the game toward an equilibrium of the information-dominant actor’s choosing. Again, in the real world, the level of capability of each actor is likely to have a decisive effect in such circumstances—an actor may understand how to manipulate an adversary’s choice structure, but find himself unable to do so. From the perspective of policy, however, the United States has the material advantage to effect changes based on superior information. Such changes may be for improvements (for example, efficiency gains) in existing policy, or for the development of new policy in pursuit of the same objectives.

The argument here sounds largely academic, but in fact our principles are derived from historical interpretation. *Seeking an information advantage to make the most of the*



World Food Programme (Gerald Bourke)

Humanitarian consequences of economic sanctions against countries might be prevented by information targeting of individuals



Airmen at Eielson Air Force Base use surveillance cameras and monitoring stations to maintain security

U.S. Air Force (Eric T. Shelley)

material context has been a linchpin of grand strategy for centuries. A prominent example is U.S. policy to dissolve European empires following World War II. As described by Peter Clarke in his *The Last Thousand Days of the British Empire*, the United States hastened the collapse of the British Empire by withholding financial assistance needed to maintain control over British territorial holdings in India and the Middle East. This dilemma for the British, referred to as a “financial Dunkirk” by contemporaries, was one that the United States created years before it happened by locking the British into the Lend-Lease system. The case demonstrates a combination of U.S. resources (diplomatic and financial) in pursuit of a grand strategy objective (postwar U.S. predominance in the West) at the cost of British might. The United States thus drew on specific knowledge of British finances and the relationship between those finances and foreign policy interests to effect a change in British behavior.

One may begin to consider the potential for advanced surveillance systems to empower similar strategy designs in today’s international system. From the perspective of U.S. foreign policy, competition in surveillance technology suggests opportunities for the furtherance of U.S. objectives, and also threats to the United States. Tapping access to foreign information networks,

Moving Forward

The examples demonstrate that the United States has an interest in developing integrated surveillance systems for offensive and defensive purposes in pursuit of a unified grand strategy framework. The offensive capabilities affirm an efficiency gain while pursuing a unified grand strategy; combining advanced behavioral modeling with the full range of U.S. policy levers suggests the ability to better achieve U.S. foreign policy objectives. The threat from others using similar approaches against the United States, however—particularly should another become the “deepest-thinking” actor by understanding U.S. preferences and incentive structures better than we understand those of our adversaries—*necessitates* drawing on new IT in pursuit of a single-equilibrium strategy.

We advocate, therefore, a reorientation of defense planning and the intelligence process supporting it. Rather than focus on the future of the international system as something that will happen to the United States, we suggest emphasizing how the United States shapes the future. This process focuses not on the many possible ways that interactions in the international system *may* unfold, but instead directs efforts toward achieving highly specific outcomes of the many interactions in world politics. As argued above, despite the promise of new IT, scenario analysis ought not be

the United States would bring its material resources to bear to execute its strategy. The third is perhaps the most important from the perspective of long-term strategy: *evaluation and refinement* refers to using the surveillance system to continue to gather and analyze data on the target of interest. This phase begs the question: “Do we know whether or not our strategy is working? If not, why not?” The strategists may then improve the surveillance system, such as by developing new or tweaking old sensors and algorithms, to constantly deploy incremental improvements that increase the scope (complexity) of social modeling. If the data flow is effective, the strategists may then understand whether or not the strategy itself is working well, and seek to refine the link between material capabilities and actions.

The introduction to this article expressed the idea of a “funneling effect” on political outcomes should the United States reorient toward a single-equilibrium strategy. We conclude by further expanding this concept. In its pursuit of a particular state of the international system or subsystems, the massive coercive authority of the United States may diminish the choices available to foreign powers and thus limit the possible futures that may come to pass. In our proposed approach, planners constantly engage in scenario analysis to consider how the trajectory of the system of interest may diverge, and how the United States may adjust its policies to return the system to a desired path. One may think of this approach as a series of “course corrections” while trying to navigate complex social “waters.” The better the United States understands its political “waters”—the output of persistent surveillance systems—the more granular it may make its “corrections”—influence operations via information networks or other tools of national power.

In sum, our ambitions are, first, to convince policymakers to get back into the game of setting explicit goals for social system outcomes; and, second, to demonstrate that the thoughtful use of IT may help the United States to achieve and preserve those desired outcomes by shaping the decisionmaking structures of adversaries and by averting crises that threaten a preferred equilibrium. The technologies empowering such a strategy process will continue to evolve, and to maximize the potential gains from revolutionary advances in IT—and minimize the gains of potential adversaries—the United States should begin planning for, and investing resources in, these breakthroughs. **JFQ**

one could learn how, given a particular cultural context, to improve the prospects of citizens “buying in” to a new regime

combined with long-range persistent surveillance, could yield unique intelligence collection systems that allow the United States to model general tendencies of population preferences in foreign countries. Such information could be useful in executing such tasks as promoting democratization (for example, one could learn how, given a particular cultural context, to improve the prospects of citizens “buying in” to a new regime) and punishing rogue states (for instance, better understanding how to target sanctions). On the other hand, one may view a parallel between the present U.S. position and that of the British at the end of World War II; massive U.S. debt, which continues to accumulate, puts foreign powers in a position to potentially manipulate U.S. preferences (and, consequently, behavior) in the long run.

discarded. Instead, scenario planning should be viewed as a critical component of perpetuating U.S. strategy. This is particularly the case as the foreseeable advancements in surveillance do not make information perfect, merely much better, with strong advantages for the strategic actor capable of “seeing” the most.

A framework for single-equilibrium strategy would consist of three primary phases: planning, implementation, and evaluation and refinement. The *planning* stage would be to identify the strategic goal, design information-gathering systems in pursuit of that goal, and use initial data collection to inform a plan to align material capabilities to attain the goal. It is during this phase that scenario planning is critical; strategists consider how, once implemented, the strategy may fail due to (or be rendered less effective by) perturbations in the social system. The *implementation* phase is self-explanatory;