

Intelligence Reform since 9/11: An Organizational Economics Perspective

Luis Garicano and Richard A. Posner¹

1. Introduction

In this paper we use insights from organizational economics to analyze the principal organizational issues raised by the nation's recent intelligence failures. We seek both to clarify the causes of these failures and, by our case studies of them, to introduce the literature of organization economics to a wider audience of economists.

The collection and analysis of intelligence are not activities limited to government intelligence agencies. Oil companies collect "intelligence" concerning the likely location of undiscovered oil fields, manufacturers of consumer products collect intelligence concerning competitors' plans and prices and consumers' tastes and intentions, and investment banks collect intelligence about political and economic conditions in foreign countries and frauds, scams, and bankruptcy risks. But the activities of government intelligence agencies differ importantly from these private-sector examples in being oriented toward protecting national security and hence lacking market measures of success.

Our analysis draws on two scholarly literatures. The first concerns the design of incentives in circumstances in which information is severely limited, though only the part of this literature that does not involve contracts that create financial incentives is relevant.² The second studies how organization affects information flows between agents:³ organizations enable individuals to circumvent the constraints of bounded rationality (innately limited cognition), so that more information can be used in a decision than an individual could process, though the increased information processing comes at a cost in terms both of delay and of loss of quality as the information moves up, down, and across the organizational hierarchy.⁴

Drawing on both literatures, we focus on three key problems that have been flagged by the high-level commissions that have studied the recent intelligence failures (National Commission on Terrorist Attacks Upon the United States [hereafter 9/11 Commission], 2004; Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction [hereafter WMD Commission], 2005). They are failures of

¹ Garicano is an Associate Professor of Economics and Strategy at the University of Chicago Graduate School of Business. Posner is a Judge of the U.S. Court of Appeals for the Seventh Circuit and a Senior Lecturer at the University of Chicago Law School. We thank Meghan Maloney for excellent research assistance; Karen Bernhardt, Dennis Carlton, Jacques Crémer, Robert Gibbons, Andrea Prat, and Luis Rayo for helpful comments on a previous draft; and Kevin M. Murphy and Jesse Shapiro for helpful discussions. For an earlier analysis of some of the issues addressed in this paper, see Posner (2005, chs. 5–6).

² For overviews of the entire literature see Gibbons (1998) and Prendergast (1999). An excellent textbook treatment is Bolton and Dewatripont (2005).

³ We use the term "agents" to designate not intelligence agents as such, but merely the individual members of any organization.

⁴ This literature builds on Arrow (1974).

analysis that result from “herding” and other deficiencies in analyzing the information in hand, illustrated by the virtually unanimous but thoroughly mistaken conclusion of the intelligence community that Iraq had weapons of mass destruction; failures to share intelligence data and analyses both laterally and vertically within the intelligence community; and the notably poor performance of the FBI as a domestic intelligence service. We discuss the three problems, and offer tentative solutions, after a brief sketch of the existing organization of intelligence in the United States.

2. The Intelligence System

Before the Intelligence Reform Act of 2004, and dating back to 1947 with relatively few changes in between, the U.S. intelligence system consisted of 15 separate agencies, including several owned by the Defense Department and dedicated to obtaining intelligence data through technical means, such as spy satellites. An official appointed by the President and called the Director of Central Intelligence (DCI) was responsible for coordinating all the agencies concerned with foreign intelligence. The agencies had (and have) overlapping functions; for example, the technical intelligence facilities owned by the Defense Department are also usable and used to collect political and other nonmilitary intelligence data, importantly including data concerning the activities of terrorists.

The Director of Central Intelligence doubled as the Director of the Central Intelligence Agency. The CIA’s centrality lay in the fact that it controlled almost all the “humint” (human intelligence—i.e., spies) and also employed most of the intelligence analysts. It collected intelligence data both on its own (as from spies run by the agency’s case officers) and from the technical and other intelligence agencies, analyzed the data, and provided the results of the analysis to the President and other high officials. Neither the CIA nor the DCI had responsibility for domestic intelligence, however; that was the responsibility mainly of the FBI, though the FBI was primarily a criminal investigation agency.

Dissatisfaction with the performance of the intelligence agencies in failing to predict the 9/11 attacks resulted in a number of changes made by the 2004 Intelligence Reform Act. Among them, the positions of DCI and Director of the CIA (DCIA) were split, and the DCI was renamed the DNI (Director of National Intelligence) and given limited budgetary, planning, and policy responsibilities, along with a broadened coordination authority that includes domestic intelligence. But he was not given command or operational authority. His relation to the DCIA is unclear, but President Bush has said that the DNI will be the President’s principal intelligence adviser, and not merely a coordinator. The Act is complex and ambiguous, so the precise limits of the DNI’s authority remain to be worked out.

A long-standing proposal that the Act did *not* embrace is to create a domestic intelligence agency separate from the FBI, on the model of the United Kingdom’s MI5 and similar agencies in other nations that have a long history of combating terrorism. MI5 is a pure intelligence agency in the sense of having no authority to arrest or to assist in criminal

prosecutions;⁵ in contrast, the FBI is a hybrid, combining general responsibility for investigating all federal crimes with a national-security intelligence function.

The WMD Commission (2005) has recommended some additional structural changes, a few of which we mention in passing. We rely on that commission's report mainly for the light it casts on the three problems that we discuss.

3. Information Analysis

A remarkable aspect of the intelligence failure in the run up to the Iraqi war was that the mistaken conclusion that Saddam Hussein had restarted his weapons of mass destruction (WMD) programs and in fact possessed WMD capabilities was shared by the entire intelligence system of the United States (indeed of the world). In the 12 years preceding the invasion, "the Intelligence Community did not produce a single analytical product that examined the possibility that Saddam Hussein's desire to escape sanctions... would cause him to destroy his WMD" (WMD Commission, 2005, pp. 155–156). We shall examine several factors that organizational economics suggests may have contributed to this unwarranted unanimity of opinion.

3.1. Corroboration versus repetition: the herding problem. A key aspect of an organizational structure is the information network—who talks to whom and thus who passes information to whom. The stages by which a particular piece of information moves from its origin to the point at which it is combined with other information for purposes of analysis are often unknown to the analyst. Yet they are the key to the reliability of the information. Analysts and consumers of intelligence therefore often require several pieces of information that confirm each other before they will believe a particular claim. But without knowledge of the structure of the network through which the intelligence has flowed, it is impossible to know how independent those confirmatory pieces of evidence really are.⁶ In concluding before the invasion of Iraq that Saddam Hussein had weapons of mass destruction, the intelligence agencies relied heavily on information supplied by Iraqi exiles, some of whose reports came through the Iraqi National Congress (INC). These reports contained similar findings, and this corroboration made them persuasive. It has since been learned, however, that rather than reflecting independent sources of information, the reports probably originated from a single source—the Iraqi National Congress itself (Reiff, 2003; Dwyer, 2004; Jehl, 2003; Isikoff and Hosenball, 2004). Another example of multiple single-source reports concerns the issue of Iraq's possession of biological WMD; all the data, except for two corroborating reports (one from INC sources), originated from "Curveball," an agent who claimed to have worked in Iraq's bioweapons program. The problem is general: "analysts were often unable to determine whether a series of raw human intelligence reporting came from the same source. For most reporting, there is currently no way to determine from the face of the CIA analyst's report whether a

⁵ Actually, it has some security functions in addition to conducting intelligence, but these are irrelevant to our analysis.

⁶ Recommendations for a job candidate illustrate this problem. Should two recommendations be taken as independent, or might one have been based on the other, in which event the two may constitute only one data point?

series of reports represents one source reporting similar information several times or several different sources independently providing the same information” (WMD Commission, 2005, p. 178).

The problem can exist even in the presence of complete rationality. When agents are ignorant of the details of the network through which information flows to them, the rational reaction may be to accept information as valid and pass it on, ignoring one’s own contradictory information. The result is information “herding” (Banerjee, 1992, Welch, 1992). Imagine a set of individuals who can observe each other and wish to decide between going to restaurant *A* and restaurant *B*. The prior belief concerning the two restaurants is that they are equally good. All the individuals have some information (assume it consists of the observation of an independent identically distributed random variable), and they know that the others do as well. Suppose individual 1 goes to restaurant *A*. Individual 2 now has two pieces of information: his own, and the fact that individual 1 went to *A*. Individual 2 may choose to go to *A* if his information agrees, but he may decide to go even if it disagrees, because 1 went. Now individual 3 has three pieces of information—the action of 1 and 2 and his own opinion. If both 1 and 2 went to *A*, 3 may well decide to do so as well because there is now quite a bit of evidence suggesting that it is a better restaurant. The herd, or information cascade, follows (Banerjee, 1992). It follows optimally, because each individual is doing the best he can given the information he has. Yet as a result of herding, everyone may be wrong, as in the WMD case, and, right or wrong, a consensus may be less epistemically robust than the sheer number of adherents to it suggests.⁷

The Curveball case differs slightly from the standard form of herding studied in the literature in which agents observe each other’s actions (such as going or not going to Restaurant *A*) and infer from those actions (imperfectly) the information that each possesses. In the Curveball case, the members of the information network, and their consumers, receive assessments that they are unsure how to evaluate because all the agents may be relying on the same ultimate source without knowing it. Yet it is rational, as in the standard herding case, to rely on the information even if it contradicts information possessed by the recipient that is of equal value, since he has only one data point while the information that he is receiving from multiple sources is at least one data point and possibly many.⁸

Because spies are unreliable at best, alternative sources of reliable information difficult to come by, and the intelligence target likely to engage in deception, herding can be prevented only if intelligence agencies can trace different pieces of intelligence to their ultimate source *and* if the trace remains attached to the information as it moves through the system. The difficulty lies in keeping the identity of the source secret, but technology may offer a solution. An encrypted tag could be assigned to a source. Different agents

⁷ The literature on “rational” herding is large, and the above brief summary cannot do justice to it. For an overview, see Bikhchandani, Hirshleifer and Welch (1998). Another literature (Scharfstein and Stein, 1990) studies how herding can result from the effort of agents to distribute blame by signing on to the office consensus. The two types of herding are connected in Ottaviani and Sorenson (2000). See also Chevalier and Ellison (1999) and Lamont (2002) for empirical evidence

⁸ This effect is amplified if agents systematically underestimate the likelihood of repetition (De Marzo, Vayanos, Zwiebel, 2003).

with different levels of security clearance would be able to decipher different amounts of the tagged information. For example, only a handful of officers with a very high clearance level would be permitted to learn the name of the source, but officers with lower clearances might be permitted to learn whether other intelligence came from the same source and what region or time frame the intelligence had originated in.

3.2. Competitive intelligence: false positives versus false negatives. Centralizing the evaluation of ideas (e.g., intelligence analysis as distinct from the collection of the raw intelligence data) eliminates more unconventional ideas and thus ultimately generates a more homogeneous product. There are fewer gross errors, but at the same time fewer interesting possibilities are ultimately considered. Which type of mistake is more harmful may depend on the organization's information environment (Sah and Stiglitz, 1986). When the environment is unstable, the organization should be so designed as to maximize the likelihood that many fresh new ideas will be produced, even if a number of them prove in the end to be unsound; for without a diversity of new ideas the organization will be unable to adapt to a changing environment. When the environment is stable, so that perfection in adapting to it is a realistic objective achievable by merely incremental adjustments, the organization should be designed with many filters so that the errors that are made are of the type "a few good ideas were not tried out" rather than "we went ahead with some terrible ideas and damaged our franchise."⁹

Filtering implies a centralized structure with layers of hierarchy whereby experienced supervisors strain out the questionable ideas originating below. A loosely knit, decentralized structure of multiple agencies is likely to be preferable in intelligence because more ideas will percolate to the top since there will be fewer layers to intercept them. Most intelligence leads and clues lead nowhere, so that the few accurate ones are scarce and therefore valuable; generally, when good ideas are scarce, a decentralized structure is preferable as more ideas will get through the filters (Sah and Stiglitz, 1986).¹⁰ Moreover, though with an important qualification discussed below, the cost of false negatives (not pursuing a lead and so failing to avert a terrorist act) is likely to be higher than the cost of false positives (pursuing a lead that turned out to be a false alarm). A model to be avoided, therefore, is the highly centralized organization of Israeli military intelligence before the 1973 surprise attack by Syria and Egypt that began the Yom Kippur war. In deciding whether there would be an attack, the Israeli Cabinet considered only one opinion, that of the chief of military intelligence. The Agranat Commission, which studied the causes of the failure to predict the surprise attack, recommended that the intelligence system be reorganized to "ensure pluralism in the various types of intelligence evaluation" ("Israel: What Went Wrong on October 6?" 1974). Similarly, in the months preceding the Iraqi war a number of low-level officers in the CIA's Directorate of Operations expressed

⁹ 3M is an example of an organization that has been successful in coming up with new ideas. The firm is decentralized and provides its employees with ample discretion. It has a culture of forgiving failure and allowing individuals to buck senior management in pursuit of what may look like unrealizable ideas (Bartlett and Mohammed, 1995). By contrast, at Procter and Gamble, which operates in a more stable product environment, new product proposals go through 40 to 50 revisions until they reach the CEO for a decision (Herbold, 2002, p. 74).

¹⁰ The analysis of this point is complex but the point seems to hold in general. Sah and Stiglitz (1986, p. 725); see also Stiglitz (2001, p. 511.)

doubts about the veracity of Curveball's information. But their superiors disagreed and (apparently) presented a filtered, unified point of view that failed to reveal the diversity of opinion at the lower levels (WMD Commission, 2005, p. 94).

But analysis is complicated when the *dynamic* aspect of the tradeoff between false positives and false negatives is taken into account. Mistaken warnings of an imminent attack lull decision makers into ignoring future warnings that may be accurate—the “boy who cried wolf” phenomenon. To minimize this cost of false alarms, the threshold for warning must be raised. This places a limit on a decentralized structure, since decentralization increases the number of false positives. The problem is particularly acute in an environment in which agents are “gun shy” because of past failures. The natural reaction to such failures is to provide too little filtering of warnings and as a result sound too many alarms. This is in fact a feature of the post-9/11 environment: “The channels conveying terrorism intelligence are clogged with trivia. One reason for this unnecessary detail is that passing information ‘up the chain’ provides bureaucratic cover against later accusations that the data was not taken seriously. As one official complained, this behavior is caused by bureaucracies that are ‘preparing for the next 9/11 Commission instead of preparing for the next 9/11’” (WMD Commission, 2005, p. 422).

3.3. Lock-in effects. As Arrow (1974) first pointed out, a design that is optimal for information processing in a given situation is unlikely to be optimal when things change, and yet agents will be reluctant to change the design and their reluctance may be rational. For once a code and a set of information channels are created, a sunk investment has been made that will constrain the organization's reaction to a new environment.¹¹ The lock-in phenomenon is thus a consequence of organizational “culture,” the complete set of codes, understandings, firm customs, perspectives, biases, etc. that substitute for explicit contracting, bargaining, enforcing, and monitoring in a setting characterized by uncertainty (Kreps, 1990).¹² Individuals interacting with each other need to know what to expect, and the organizational culture tells them whether a given performance is acceptable. As in the case of social norms, changing an organization's culture is difficult because the culture is embedded in a network of complicated, informal personal interactions and expectations. The WMD Commission (2005, ch. 10) gives many examples of how the organizational culture of the FBI has thwarted the efforts of the FBI's director to recast the Bureau as an intelligence service rather than merely a criminal investigation agency.

Lock-in effects help explain the difficulty that intelligence agencies had adapting to the end of the Cold War. As noted by the WMD Commission (2005, p. 4), the organizational structure of the intelligence community had evolved to deal with a focused threat, that of

¹¹ . Thus in their analysis of the photolithographic equipment industry, Henderson and Clark (1990) find that firms successful in one generation often had an organizational architecture that made it difficult for them to recognize changes in demand and technology that would dominate the market in the next generation. Such constraints, it should be noted, can also operate at an individual, rather than organizational, level. Mullainathan (2000) argues that the mind assigns a particular observation to a previously formed category; which limits the possibility of updating on the basis of new information. Only major shocks will induce a change in the categorization scheme.

¹² See Crémer (1995) for an alternative view of culture as a stock of shared specific human capital within the organization, composed of knowledge of facts, rules of behavior, and a code.

the Soviet Union, with its concentrated nuclear and conventional forces. The current threat environment includes dozens of state and nonstate entities that could strike the United States—and in some cases with weapons that are difficult to detect by means of the intelligence apparatus that had been designed to meet the earlier challenge. This is the “failure of imagination” emphasized by the 9/11 Commission’s report (9/11 Commission, 2004, pp. 339–348).

Reducing the effects of lock-in is another reason to prefer multiple competing intelligence agencies. Like competing firms in the market that outmaneuver those that “don’t get it,” individual agency biases caused by lock-in of past information structures will distort the aggregate intelligence product less if there is competition; and there will be a greater chance that at least some of the organizations will prove well adapted to an altered environment, should the current environment change. Some of the cultural differences in the intelligence community are dramatic: think of the very different perspectives of the Defense Intelligence Agency, reflecting the outlook of the military; the CIA, with its tradition of clandestine activity; and the State Department, with its traditional hostility to military adventures and clandestine activity (which frequently provokes diplomatic incidents).

3.4. Incentive problems

When supervisors are trying to evaluate the quality of subordinates rather than just their effort, subordinates may be able to manipulate the supervisors’ opinions of quality (Holmstrom, 1982).

3.4.1. “Yes men.” An agent who knows what his manager thinks may be inclined to bias his report toward agreement with the manager’s views; the manager is likely to think more highly of him if his report coincides with the manager’s preconceptions. In other words, agents may become “yes men” (Prendergast, 1993). Before the Iraqi war, intelligence officers knew that the policymakers—their ultimate superiors, believed that Iraq possessed weapons of mass destruction. The Department of Energy, however, believed that a key datum that had led other agencies to conclude that Iraq was trying to develop nuclear weapons (aluminum tubes that could be used in the manufacture of such weapons) was misunderstood by the other agencies. The Department relied on a separate piece of evidence, doubted by the other agencies, that Iraq was seeking a nuclear weapons program.¹³ The WMD Commission believed that the Department of Energy was reluctant to buck the consensus view of Iraq’s WMD capability.¹⁴

This phenomenon need have nothing to do with overt “politicization” of intelligence, that is, pressure from superiors. All that is required is for the agent to know the supervisor’s

¹³ “The Department of Energy (DOE) assessed that the tubes ‘were not well suited for a centrifuge application’ and were more likely intended for use in Iraq’s Nasser 81 millimeter Multiple Rocket Launcher (MRL) program” (WMD Commission, 2005, p. 56).

¹⁴ The WMD Commission found that the Department of Energy intelligence analyst who participated in the coordination meetings for the NIE “conceded to this Commission that ‘the DOE didn’t want to come out before the war and say [Iraq] wasn’t reconstituting’” (2005, p. 75).

preconceptions. Intelligence managers are said to have rewarded judgments that reflected the consensus view that Iraq had WMD programs, and penalized judgments that did not (WMD Commission, 2005, p. 191).

Minimizing the “yes men” phenomenon requires that governance be decentralized. If there is an all-powerful intelligence CEO, his intuitions, perspective, etc. will exercise a magnetic force over the advice he receives from below. If, however, there are multiple, more or less independent agencies, diverse views will survive even if, within each agency, subordinates hesitate to buck their superiors.

3.4.2. Failure to update on the basis of new information. Prendergast and Stole (1996) argue that as agents become more experienced, they may update too little in the face of new information. Frequent updating may lead others to judge the updater to have unreliable information, whereas a manager who sticks to his opinion signals that he’s had good information all along; in effect he is herding with himself. The tendency is greater with more experienced managers because they have more of a track record, meaning that any bold departure they make is more likely to contradict a position they’ve taken in the past. Only if they can credibly claim that their information is new can they change their minds without losing reputation.

The need to be seen to be consistent plays a role even at the agency level. The CIA defended the Curveball intelligence long after it was discredited, fearing how an acknowledgment of error would look to senior management at the CIA and to policymakers (WMD Commission, 2005, p. 107).

Avoiding these biases may require shifting agency managers regularly, so that they are not bound to previous forecasts and analysis. But this is not a panacea; as Prendergast and Stole (1996) point out, new managers may take positions that are *too* bold, in order to signal that they must be well informed.

3.4.3. “Consumer” preference. Even in a setting of competitive intelligence, if the ultimate consumers of intelligence—the President and other high-level policymakers—are homogenous in their preconceptions, intelligence officers will be tempted to anticipate those preconceptions and shape intelligence advice to conform to them. In an analysis of the market for news, Mullainathan and Shleifer (forthcoming) show that reader (i.e., consumer) heterogeneity plays a more important role in enhancing diversity of opinion in the media than competition does. When readers have homogeneous biases, competition does not counteract bias in the media, because the media cater to the biases and hence tastes of their consumers. When readers’ biases are heterogeneous, competition generates heterogeneous media biases (the media chasing their consumers—conservative readers read conservative media, liberals read liberal media), which will tend to offset each other if readers are willing to read multiple sources (few are).

“Consumer sovereignty” is also the rule in the intelligence business. The influence and budgets of the intelligence agencies depend ultimately on the President, members of Congress, and other high-level politicians and policymakers, who may be angered by in-

telligence analysis that contradicts their priors. Competing producers of news, here intelligence agencies, must thus be complemented by diverse consumers, here policymakers, for unbiased intelligence to be produced.

4. Information Sharing

Lack of information sharing within agencies (horizontally and vertically), between different agencies, and between federal, state, and local government levels has been blamed for the intelligence failures that took place in relation to both 9/11 and Iraq. The CIA did not pass on the names of suspected terrorists to the Federal Aviation Administration before 9/11; the memo from the FBI's Phoenix field office requesting investigation of flying schools was never followed up;¹⁵ the reports questioning Curveball's credibility were not disseminated widely enough within the intelligence community and to policymakers (WMD Commission, 2005, p. 430); and so on. In what follows, we consider from both an information-processing and an incentive perspective why information is not shared.

4.1. Information structure

Arrow (1974) suggested two features that enable an organization to acquire more information than an individual could. They are specialized codes and hierarchy.

4.1.1. Codes. An organizational code is a technical language that agents use to communicate among themselves. It consists of both an ability to share digitized data by standardizing formats, commands, and protocols and a common vocabulary ("terrorism analysts even used a different vocabulary [from traditional WMD analysts] to describe unconventional weapons capabilities, using the term 'CBRN'—chemical, biological, radiological, and nuclear-weapons programs instead of 'WMD' programs" (WMD Commission, 2005, p. 275)). Crémer, Garicano, and Prat (2005) study the tradeoff between a specialized code, which improves communication within the organization but tends to isolate it from the outside, and a less specialized code that facilitates coordination among organizations but impedes communication within the organization. A competing consideration, however, is that if different organizations with overlapping functions (such as the different intelligence agencies) have different codes, some of the organizations may prove better adapted to a change in the environment, as in natural selection. (This is similar to our earlier point about the value of diversity in combating lock-in.) Thus the short-term advantage of code uniformity in facilitating communication across agencies has to be traded off against the long-term advantage of code diversity in facilitating the system's adaptation to change.

4.1.2. Hierarchy. Organizational hierarchy enables the aggregation of information. Each agent processes a bit of information, summarizes what matters or combines it with other information, and passes the summary up to the next level.¹⁶ Hierarchy also enables experts' knowledge to be reserved for situations in which their knowledge is especially

¹⁵ See Chapter 8 of the 9/11 Commission report for an account of the miscommunications and leads not followed.

¹⁶ For an early review of formal models of hierarchies as information aggregators, see Radner (1992).

valuable. A bit of information from the field is, if routine, dealt with by a field officer. If the bit is unusual, it goes up the organizational ladder to a more knowledgeable officer who decides what action to take. Information that is truly exceptional continues up to the top of the hierarchy. This “management by exception” allows for the optimal matching of problems with expertise.¹⁷

A downside of information aggregation in hierarchy is that information tends to be garbled in successive transmissions and some has to be discarded deliberately in order to avoid overloading the top echelon, where the information flows from the different divisions of the organization come together (Williamson, 1967). Thus “criticism of Curveball grew less pointed when expressed in writing and as the issue rose through the CIA’s chain of command” (WMD Commission, 2005, p. 105). Similarly, the presidential daily briefing suppresses most of the nuances of the information being presented: “policymakers are sometimes surprised to find that longer, in-depth intelligence reporting provides a different view from that covered by the PDB” (WMD Commission, 2005, p. 420).

4.1.3. Improving information channels: the role of information technology. Both the WMD and 9/11 commissions noted that information technology can improve the sharing of information. The scholarly literature points out that organizational change is often necessary, however, to enable organizations to profit maximally from that technology. It has been found that organizations that take full advantage of IT become more decentralized (more decisions are made by frontline employees because electronic databases and personal computers put more knowledge at their fingertips), use higher-skilled labor, and have fewer hierarchical layers and therefore larger spans of control for each manager.¹⁸ IT also alters the tradeoff between codes and hierarchies (Crémer, Garicano, and Prat, 2005). An important function of managers is to translate among subordinates who use different codes. With digitization enabling specialized codes to be integrated into common codes, the role of managers in facilitating information transfer declines and organizations tend therefore to flatten out.¹⁹

A complication in the intelligence arena is that the need to maintain secrecy, for example of spies’ names and of the extent to which the intelligence system is able to eavesdrop on enemies, limits horizontal sharing. Often before intelligence can be shared across agencies or with policymaking officials it must be “scrubbed” to eliminate clues to secret information. This may require that intelligence travel some distance up the hierarchy of the

¹⁷ A formal analysis of how hierarchies allow the knowledge of experts to be conserved, and of the conditions under which such “management by exception” hierarchies is optimal, is in Garicano (2000).

¹⁸ See Brynjolfsson and Hitt (2000); Bresnahan, Brynjolfsson and Hitt (2002); and Caroli and Van Reenen (2001) for evidence of the relation between IT and decentralization; and Rajan and Wulf (2003) for evidence on the impact on hierarchical spans and layers. Communications technology, as opposed to “information technology” in the sense of digitization, manipulation, etc., of data, increases centralization by enabling more decisions to be made by higher-level decisionmakers. See Garicano (2000).

¹⁹ Examples are the design of the B-2 bomber—engineers were allowed to communicate across firm boundaries without relying on senior management (Argyres, 1999)—and the unification of human resource management and financial database management inside Microsoft (Herbold, 2002), which enabled a substantial increase in horizontal communication. Crémer, Garicano, and Prat (2005) review the evidence in light of the theory.

agency in which it was obtained before it can be shared with another agency (or another unit of the same agency). Hence a national-security intelligence agency can be expected to have a more hierarchical management structure than its private counterpart.

4.2. Incentives for sharing.

Besides having the right structure of links, codes, etc. in the organization, incentives must be created for individuals to share information with others so that analysis rests on as broad a basis in data as possible. This need presents unique challenges in comparison to the production and distribution of other goods. It is difficult to know what others may know or to make someone share his knowledge against his will and the value of knowledge to the holder is often destroyed once the knowledge is shared.

4.2.1. Sabotage and influence activities. The creation of financial incentives for employees in public agencies takes the form not of gearing salary or bonus to the actual output produced, as in the private sector, because that output cannot be monetized, but rather of gearing it to supervisors' assessments that are based in turn on observing the agents' performance. Thus employees compete against each other for promotion, and while such tournaments can have good incentive effects (Lazear and Rosen, 1981), they can also have bad ones. Agents may try to sabotage each other (Lazear, 1989) by concealing information or providing false information. Or they may squander resources on manipulating the perception of their performance by superiors or otherwise gaining the favor of those superiors—what the literature calls “influence activities” (Milgrom and Roberts, 1988, 1990). An example is the presidential daily briefing, which has become the primary platform by which intelligence agencies seek to advertise their products in competition with each other:²⁰ “The daily reports seemed to be ‘selling’ intelligence—in order to keep its customers, or at least the First Customer, interested” (WMD Commission, 2005, p. 14)

4.2.2. Turf wars. An extreme version of influence activities is the turf war. An agent (or agency) can allocate resources either to productive activities or to gaining appropriations, personnel, or other advantages at the expense of another agency. Since the probability of victory in the “war” depends on the resources devoted to it, the less productive agency will devote more resources to winning and thus will tend to win, though by definition it is the less efficient producer (Skaperdas, 1992).

Turf wars are highlighted in every report on intelligence and other governmental failure. Examples include (1) the FBI versus the CIA—a turf war at the heart of the pre-9/11 failure to track the foreign terrorists as they entered the US (9/11 Commission, 2004, p. 263), and continuing today; for example, officials at the CIA's Counterterrorism Center claim that “they have difficulty tracking and obtaining information about terrorist cases after they hand them off to the FBI” (WMD Commission, 2005, p. 469); (2) the CIA versus the intelligence agencies controlled by the Defense Department (WMD Commission, 2005, p. 332); and (3) the CIA's Counterterrorism Center versus the newly created National Counterterrorism Center outside the CIA. Besides wasting resources, turf wars in

²⁰ The tournament method of promotion will also, however, tend to make agents disperse rather than herd, for only by dispersing can the individual distinguish himself (see Zwiebel, 1995).

the intelligence community retard the sharing of information because sharing confers a benefit on the rival agency.

Turf wars can be exacerbated by conflicts of interest, such as the conflict of interest that the Department of Defense has as both the owner, and one of several customers, of the principal technical intelligence services, such as the National Security Agency, which intercepts electronic communications. In a market setting, such vertical integration is common and not problematic. A steel company, for example, may manufacture both sheet steel and finished steel products which it sells in competition with fabricators who buy their sheet steel from it. Competition constrains the steel company to sell to its fabricator competitors at a competitive price. The Defense Department, however, cannot charge for the use of its technical intelligence facilities by the CIA or other nondefense intelligence agencies, and so it has a disincentive to share, since sharing reduces the amount of purely military intelligence and thus encroaches on DoD turf. There is no mechanism for trading off costs and benefits of different uses of the facilities. This problem could be resolved by spinning off the technical agencies into their own, free-standing technical intelligence agency, which would have no incentive to favor military over other customers for its services.

4.2.3. Information rents. Agents may not share intelligence because they do not want to lose the rents derived from their control of the resulting knowledge. An FBI agent who has information that may lead in the future to an arrest may realize that passing this information to another agent, or to another agency, is the right thing to do from a social perspective. But if he does this he will be dissipating some of the rents generated by his monopoly of the information—he will be unlikely to make the arrest and earn the consequent career rewards.²¹ For an agent to share information in an electronic database decreases his uniqueness and expertise, and thus his power and rents. This weakness of the incentives for sharing, as a result of “information ownership” incentives, was emphasized in the 9/11 Commission’s report and continues to plague the intelligence system. The WMD Commission’s report found that “individual departments and agencies continue to act as though they own the information they collect” (WMD Commission, 2005, p. 14).

4.2.4. Designing incentives for sharing. If agents are rewarded for sharing information, the quality of the information may suffer, in much the same way that quality suffers if case officers are rewarded according to the number of spies they recruit. What is needed is a method of determining how valuable the information is to the recipients. One way to do this is, by means of the “tag” system suggested in section 3.1, to count the number of times the information is cited in analytical reports. Analogies are the use of number of “hits” on a blog site to determine advertising rates for the site and the use of number of citations to an academic’s published scholarship as a basis for tenure decisions. The danger is that agents would use their contacts and influence to obtain additional citations.²²

²¹ See Garicano and Santos (2004) for an analysis of this problem in a profit-making environment.

²² This is an instance of the “you get what you pay for” problem that we discuss at greater length in section 5.2.2.

Even in the absence of strong incentives, an organization can nudge agents in the desired direction by lowering the costs of undertaking the actions that the organization wants to encourage. To lower the cost of sharing information, agents from different divisions (or agencies) with related responsibilities could be officed in close proximity, as in the case of the National Counterterrorism Center. Akerlof and Kranton (2005) and others have pointed out that an employee's sense of identification with his employer will tend to align individual and organizational incentives and thus reduce principal-agent conflict. Military organizations, and paramilitary organizations such as intelligence agencies, endeavor to create an esprit de corps that substitutes for financial incentives to good performance. In the case of intelligence, agents who care deeply about their mission may be induced to share intelligence data even if it is detrimental to their career.

5. Mission Incompatibility at the FBI

The 9/11 Commission's report identified shortcomings in the FBI's performance that made the United States highly vulnerable to terrorist attacks such as those of 9/11. Specifically, the commission found that the FBI had failed both to collect adequate intelligence data and to combine the raw, disaggregated data into accurate knowledge of terrorist threats.

In the eight years following the first World Trade Center bombing (the 1993 truck bombing), the FBI tried without success to develop an effective domestic intelligence capability. On two separate occasions, it adopted strategic plans that it failed to implement. The 9/11 Commission identified several causes of such failures: FBI headquarters encountered obstacles and resistance from the local field offices; despite pledges by Congress and the Department of Justice, additional resources were not allocated to the FBI's intelligence activities; intelligence analysts had to rely on personal relationships to obtain intelligence data from the files of local field offices, because the data in the different offices were not fed into a single database or otherwise aggregated and shared, and hence patterns of terrorist activity were not detected; the required human resources were never fully developed (9/11 Commission, 2004, pp. 76-78).

Most critics of the FBI's poor performance blame flawed management and execution. A deeper analysis suggests that the problem is not individual or collective incompetence but unsound organizational design.

5.1. Organization for crime fighting.

Historically, and especially after abuses of civil liberties by the FBI surfaced in the 1970s, and after the Cold War ended in 1989, the FBI had concentrated on investigations of ordinary (that is, unrelated to national security) federal crimes, such as bank robbery, bank fraud, mail fraud, public corruption, and drug offenses. Consistent with this focus, the organization was optimized to the law enforcement function, which involves investigation, arrest, and providing evidence for use in prosecuting criminals. The law-enforcement process is reactive. The emphasis is on catching suspected criminals and

building legal cases that attorneys in the Department of Justice can successfully prosecute. Although minimizing crime is understood to be the ultimate goal of law enforcement, the proximate goal is punishing identified criminals. Crime is presupposed and the object is not to prevent crimes from being committed but to punish the perpetrators.

Decentralized decision making and a performance measurement and reward system based upon arrests, indictments, and convictions are the key organizational attributes supporting the law enforcement function (9/11 Commission, 2004, pp. 108–109). Historically, FBI headquarters delegated decision “rights” for the initiation and conduct of cases down to the field offices. Each office “owned” the cases it initiated, and controlled all work on the case. The initiating office is called the “office of origin,” and the “office of origin mentality” is central to the FBI’s organizational culture. The office of origin would control even a case involving criminal activity that overlapped the jurisdiction of other local offices.

So decentralized a system is likely to work well in crime fighting as long as most crime is local or regional rather than national or international; and most federal crime *is* local or regional. The knowledge needed in crime fighting is mainly of the type that Hayek (1945) called “knowledge of time and place” and such knowledge normally is local. Moreover, as Aghion and Tirole (1997) have noted, delegation functions as a commitment by the center not to intervene.²³ This commitment motivates FBI agents to invest in new local knowledge, confident that agents from headquarters are unlikely to swoop down and make the arrests and build the prosecutions. Agents are more likely to collect information assiduously when they know they can decide how to use it; their ability to retain the case gives them a substitute for a property right. Decentralization has the further advantage of enabling yardstick competition—performance can be assessed better by comparing the performance of the different offices or divisions of the organization.²⁴

A decentralized system has costs as well as benefits, however. Once agents can appeal to local knowledge as the basis of their action, the center, which by definition lacks that knowledge, cannot monitor their decisions effectively. With this loss of control, agents may be able to take actions that advance their careers but do not benefit the organization as a whole. These problems are not acute when all the agency does is criminal investigation. The outputs of FBI agents engaged in criminal investigation—such outputs as number of arrests, prosecutions, convictions and affirmances of convictions, length of sentences, and amount of property recovered—are quantitative and therefore easy to measure and monitor, and, being difficult to (legally) manipulate, are credible—baseless arrests undertaken for career advancement can be detected by observing the low likelihood of convictions resulting from such arrest. For these reasons, the cost of being unable to monitor inputs effectively may not be high. As Baker (1992) notes, however, a key factor

²³ Of course, the center cannot actually ‘commit’ not to intervene, as it can always take the power back from the field offices, but rather the promise not to intervene is supported by a relational contract between the local offices and the center which the center respects to preserve local incentives (Baker, Gibbons and Murphy, 1999).

²⁴ More precisely, a decentralized structure does not affect the incentives of the lowest-level employees, nor those of the top managers, but improves the incentives of the middle managers by subjecting them to yardstick competition. Maskin, Qian, and Xu (2000) analyze this phenomenon with an interesting application to the difference between the Soviet and Chinese communist systems.

in evaluating the adequacy of a performance measure is the correlation between the effect of the agents' actions on the measure and the effect on what the principal actually desires. But that condition is satisfied; if agents focus on increasing arrests and convictions, this will tend to reduce the crime rate. So decentralization of FBI criminal investigations is not problematic, and for the additional reason that a frequent cost of decentralization—lack of coordination—is of limited importance if most crime is local, so that coordination between field offices is not required; by the same token, the center's knowledge is less likely to be useful than if crime were national.

5.2. Organizing for counterterrorism

Collecting, analyzing, and integrating intelligence differ greatly from crime fighting. Intelligence tries to detect plots, threats, etc. *before* they reach the level at which they are prosecutable crimes. It aims at long-term penetration of suspect groups, and thus does not worry about collecting evidence of criminal guilt—it assumes that once subversive activity is discovered, it can be disrupted by disclosure, disinformation, bribery, blackmail, deportation, etc. It actually disfavors criminal prosecutions, because they reveal to the enemy what the government knows and because, trials being public, secret information tends to be disclosed in them. Not being oriented to prosecution, an intelligence service doesn't have to worry about whether the data it collects would be admissible as evidence in a criminal trial. These characteristics greatly limit the organizational value of quantitative output measures such as arrests and convictions.

5.2.1. Initiative versus coordination. As Roberts (2004) has pointed out, organizations in which authority is decentralized encourage individual initiative at a possible sacrifice of coordination. Hierarchies, by contrast, sacrifice some initiative to ensure coordination. In particular, when individual performance incentives are strong, as Dessein, Garicano, and Gertner (2005) show, communication among agents is not credible, as any attempt to impose joint, coordinated decisions (which are not individually optimal) are met with (noncredible) messages by the agents that the costs of such coordination is too high. Thus while decentralized evaluation of ideas and decentralized, competing information collection and analysis (in the form of multiple competing agencies) may be desirable in an intelligence system for the reasons we emphasized earlier, decentralized authority over resource allocation may result in insufficient coordination. Thus, with their decentralized decision rights, the local FBI offices are motivated to exert effort but not to collaborate with one another and particularly not to acknowledge that they might be able to do without certain resources—that those resources might be employed more effectively elsewhere. Both before and after 9/11, the local offices responded to requests for counterterrorism effort by asserting the unavailability of human resources. The pre-9/11 FBI did not fail in the intelligence mission because of a lack of initiative by field offices (on the contrary, pre-9/11 agents in Phoenix and Minneapolis proved able to follow through on their hunches and obtain the relevant information (9/11 Commission, 2004, pp. 272–276)), but because of a failure to coordinate the information generated in the different offices

5.2.2. Measurability. As Holmstrom and Milgrom (1991, 1994) point out, in an environment with multiple tasks that are observable with different difficulty the establishment of clear performance criteria for the tasks that are easily measurable deflects agents' efforts from tasks that may be more valuable to those designing the incentives but also are more difficult to measure. A mere verbal declaration of commitment to intelligence is unlikely to cause FBI agents to abandon the safe and traditional career path associated with criminal investigation. The uncertainties concerning the causal linkage between intelligence activity and the prevention of terrorist acts are so great that linking career success to success in preventing terrorism acts may be tantamount to linking it to the weather: such a linking would simply add risk to the individual intelligence officer's career. Terrorism acts are few and far between and probably would be even if the nation's intelligence capability were poor, so the fact that a terrorist act is not committed cannot be confidently assumed to be an output of intelligence activity. Indeed, the earlier in the planning stage a terrorist project is interrupted, the less likely it is that, had the planning been allowed to continue, the project would actually have been carried through rather than abandoned. (Much terrorist planning turns out to be just talk.) In addition, effective intelligence depends on the performance of many agents and organizations, rather than on an individual or a small team, as in the case of most criminal investigations. A related point is that the link between individual tidbits of information and sound knowledge based on them is often unclear. Finally, the social benefits derived from intelligence activities lie not alone in information obtained, but also in simply increasing the costs of the enemy by forcing him to use more cumbersome methods of communication (or to communicate less) or to increase his efforts to prevent his organization from being penetrated. Such benefits are exceedingly hard to assess.

These measurement difficulties have made efforts to devise "objective" criteria for advancement in an intelligence service perverse. Intelligence analysts are frequently rewarded on the basis of the sheer number of finished intelligence pieces produced, which encourages them to focus on short-term issues that can be analyzed quickly (WMD Commission, p. 175.), and case officers on the basis of how many spies they recruit, which encourages them to focus on individuals who are easy to recruit and as a result unlikely to have the best access to valuable information (WMD Commission, p. 159).²⁵ A better alternative might be a system in which agents who perform adequately advance on the basis of seniority and those that do not perform adequately lose their jobs. To minimize internal conflict and encourage cooperation, including sharing of information., departures from seniority should be limited to extreme cases of incompetent or superlative performance.²⁶ A seniority-based promotion system is feasible when the employer is a

²⁵ See Gibbons (1998) and Prendergast (1999) for examples of "you pay what you get" problems in domains unrelated to intelligence. For example, Anderson, Burkhayser and Raymond (1993) and Cragg (1997) find that the Job Training Partnership Act, which keyed rewards to the number of persons reemployed, induced recruitment of persons likely to find jobs without retraining—it was easier and cheaper to find jobs for them. Another, familiar example is the substitution of quantity for quality as a consequence of the practice of paying lawyers by the number of hours they work on a case.

²⁶ Effectively, this would mean substituting for the explicit objective measures that exist now with implicit contracts based on subjective measures sustained by the parties' desire to maintain a good reputation to facilitate future advantageous transactions (Bull 1987). On the interaction between subjective and objective performance measures, see Baker, Gibbons and Murphy (1994).

monopsonist, as in the case of the military or, to a lesser extent, the intelligence services. The reason that lockstep seniority-based compensation schemes break down in the private sector—imagine a law firm in which all partners of the same seniority are paid the same share of the partnership income—is that the ablest are undercompensated and leave for a firm that doesn't have a lockstep system. But someone who wants a career in the military or in national-security intelligence cannot find an equivalent position in the private sector in which he would be paid and promoted on the basis of merit.

5.2.3. Selection. When output cannot be monitored effectively, an organization must shift its focus to monitoring the quality of the inputs (e.g., Prendergast, 2002). This suggests that the FBI will screen intelligence applicants more carefully, and they will therefore be of higher quality, than the criminal investigators whom it hires. This can be a source of tension within the organization. In addition, it suggests that the FBI will not be able to economize by using FBI agents interchangeably as criminal investigators and as intelligence officers. An offsetting consideration should be noted, however. When an agency engages in two activities, in one of which output can be monitored readily and in the other not, abler applicants may prefer the former, though the agency itself may wish to channel its ablest applicants into the activity in which performance is more difficult to assess. The able applicant expects to excel in either activity, but in the first his merit is more likely to be perceived by his superiors. The result may be a mismatch between applicants and jobs and a disparity in average quality between the two activities, but the disparity is in the opposite direction from the one suggested above. The two effects may cancel, but they need not.

Similar issues arise in other types of organization. A salesman's output is readily monitored; a lawyer's is not; yet both may work for the same organization, requiring different career tracks. But a complicating factor in the case of the FBI is the imbalance between the two functions of criminal investigation and of intelligence. More than 90 percent of FBI agents are involved in criminal investigation rather than in intelligence related to national security.²⁷ And almost all the supervisory personnel in the Bureau are career criminal investigators rather than intelligence officers. Hence the culture of the Bureau is dominated by the personnel and other practices that are optimized to criminal investigation rather than to intelligence. At the least, there should be separate career tracks for criminal investigators and intelligence officers, given the possibility of systematic differences in ability and the need for evaluating performance differently in the two jobs.

5.3. A separate domestic intelligence agency?

A more radical suggestion that our analysis supports is to remove the domestic intelligence function from the FBI and lodge it in a separate agency modeled on the United Kingdom's MI5 (the official name is the "Security Service"), which would have no criminal-investigation function. In such an agency, intelligence would not be seen as a second-class job; it would be *the* job of the agency. By having an agency specialize in the difficulty-to-measure intelligence tasks, with lower-powered incentives (i.e., avoiding quantitative measures for rewarding employees); and another agency specialized in the

²⁷ The FBI also conducts intelligence related to ordinary criminal activity.

easier-to-measure law enforcement tasks, with strong incentives, the risk of individuals' substituting effort away from the tasks that are hard to measure to those that are easy to measure is reduced (Holmstrom and Milgrom, 1991), without need to reduce incentives all around. More generally, such a separation would recognize the profoundly different organizational culture, in the sense of Kreps (1990), that is appropriate to national-security intelligence as distinct from criminal investigation; in intelligence, implicit contracts and subjective performance measures (Bull 1987) would dominate quantitative criteria, unless wholly inappropriate such criteria, such as number of intelligence reports written, were adopted in desperation.

A puzzle that we are not able to resolve is *why* the FBI wants to engage in national-security intelligence. Diversification is not a plausible explanation; and risk aversion cuts the other way, since intelligence has a bigger downside than upside. If an intelligence agency fails to predict an attack, it is blamed; but if it does predict the attack, and as a result the attack is foiled, the agency is unlikely to get commensurate credit. There are two reasons. First, the agency may conceal its successes in order to avoid providing useful information to its enemies. Second, an attack that occurs is a vivid, memorable event; an attack that does not occur is much less so. There is bound to be uncertainty whether the attack would have occurred even in the absence of discovery by the intelligence agency, and there is an imagination cost incurred in assessing the benefit of a deflected attack compared to the cost of an actual attack.

6. Conclusion

Organization economics flags a number of issues that must be considered in the design of an organizational structure. These issues arise in an acute form in the case of national-security intelligence. One cluster of issues concerns the obstacles to accurate analysis of intelligence. Partly because of secrecy concerns, the herding problem is especially acute in intelligence; there are also serious lock-in problems. To some extent the difficulty of achieving accuracy in intelligence analysis is inherent in the limited information-processing capacity of individuals and organizations. But it is also a function of incentives, for example to cater to the preconceptions of one's superiors. The second cluster of issues has to do with reluctance to share information across individuals and also across agencies. Here incentive considerations work strongly against sharing, but we have suggested a method that might tend to counteract those pressures. Finally, a clash of organizational cultures, resulting from asymmetries in the measurement of performance, suggests that it is a mistake to attempt to combine intelligence and criminal investigation in a single agency (the FBI).

The joinder of information problems with incentive problems helps explain why intelligence failures appear to be common, even when the costs of such a failure are very great, as in the examples discussed in this paper. The inherent, systemic obstacles to good intelligence performance should give policymakers pause in proposing radical solutions. Modest solutions, along the lines that we have suggested, seem more likely to produce genuine improvements.

References

Aghion, Philippe and Jean Tirole. 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy*. 105:1, pp. 1-29.

Akerlof, George A. and Rachel E. Kranton. 2005. "Identity and the Economics of Organizations." *Journal of Economic Perspectives*. 19:1, pp. 9-32.

Anderson, Kathryn, Richard Burkhauser, and Jennie Raymond. 1993. "The Effect of Creaming on Placement Rates under the Job Training Partnership Act," *Industrial and Labor Relations Review*. 46, pp. 613-24.

Argyres, Nicholas S. 1999. "The Impact of Information Technology on Coordination: Evidence from the B-2 'Stealth' Bomber." *Organization Science*. 10:2, pp. 162-180.

Arrow, Kenneth. 1974. *The Limits of Organization*. New York: Norton.

Baker, George P. 1992. "Incentive Contracts and Performance Measurement." *Journal of Political Economy*. 100:3, pp. 598-614.

Baker, George, Robert Gibbons, Kevin J. Murphy. 1994. "Subjective Performance Measures in Optimal Incentive Contracts." 109:4, pp. 1125-1156.

Baker, George, Robert Gibbons, and Kevin J. Murphy. 1999. "Informal Authority in Organizations." *Journal of Law, Economics, and Organization* 15:1, pp. 56-73.

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*. August, 107:3, pp. 797-817.

Bartlett, Christopher A. and Afroze Mohammed. 1995. "3M: Profile of an Innovating Company." Harvard Business School Case 9-395-016.

Bikhchandani, Sushil, David Hirshleifer and Ivo Welch. 1998. "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades." *Journal of Economic Perspectives*. 12:3, pp. 151-170.

Bolton, Patrick and Matthias Dewatripont. 2005. *Contract Theory*. Cambridge: MIT Press.

Bresnahan, Timothy F., Erik Brynjolfsson and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *Quarterly Journal of Economics*. February, 117:1, pp. 339-376.

Brynjolfsson, Erik and Lorin M. Hitt. 2000. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *Journal of Economic Perspectives*. 14:4, pp. 23-48.

Bull, Clive. 1987. "The Existence of Self-Enforcing Implicit Contracts." *The Quarterly Journal of Economics*, 102:1, pp. 147-160.

Caroli, Eve and John Van Reenen. 2001. "Skill-Biased Organizational Change? Evidence from a Panel of British and French Establishments." *Quarterly Journal of Economics*. November, 116:4, pp.1449-1492.

Chevalier, Judith and Glenn Ellison. 1999. "Career Concerns of Mutual Fund Managers." *Quarterly Journal of Economics*. 114:2, pp. 389-432.

The Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction. 2005. *Report of the Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction*. Washington, D.C.: Government Printing Office. Available at <http://www.wmd.gov/report/wmd_report.pdf>.

Cragg, Michael, "Performance Incentives in the Public Sector: Evidence from the Job Training Partnership Act." 1997. *Journal of Law and Economics and Organization*. 13. pp. 147-68.

Crémer, Jacques. 1995. "Corporate Culture and Shared Knowledge." *Industrial and Corporate Change*. 2, pp. 351-386.

Crémer, Jacques, Luis Garicano and Andrea Prat. 2005. "Codes in Organizations." Mimeo, University of Chicago. 2005.

De Marzo, Peter M., Dimitri Vayanos and Jeffrey Zweibel. 2003. "Persuasion bias, social influence and unidimensional opinions." *Quarterly Journal of Economics*, 118: 3, pp. 909-968.

Dessein, Wouter, Luis Garicano and Robert Gertner. 2005. "Organizing for Synergies." Mimeo, University of Chicago.

Dwyer, Jim. 2004. "Defectors' Reports on Iraq Arms Were Embellished, Exile Asserts." *New York Times* (final ed.). July 9.

Garicano, Luis. 2000. "Hierarchies and the Organization of Knowledge in Production." *Journal of Political Economy*. October, 108:5, pp. 874-904.

Garicano, Luis and Tano Santos. 2004. "Referrals." *American Economic Review*. June, 94:3, pp. 499-525.

Gibbons, Robert. 1998. "Incentives in Organizations." *Journal of Economic Perspectives*. Autumn, 12:4, pp. 115-132.

Hayek, F. A. 1945. "The Use of Knowledge in Society." *American Economic Review*. September, 35:4, pp. 519-530.

Henderson, Rebecca M. and Kim B. Clark. 1990. "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms." *Administrative Science Quarterly*. March, 35:Special Issue, pp. 9-30.

Herbold, Robert J. 2002. "Inside Microsoft: Balancing Creativity and Discipline." *Harvard Business Review*. January, 80:1, 72-79.

Holmstrom, Bengt. 1982. "Managerial Incentives Problems – A Dynamic Perspective," in *Essays in Economics and Management in Honor of Hans Wahlbeck*. Helsinki: Swedish School of Economics.

Holmstrom, Bengt and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, & Organization*. 7:Special Issue, 24-52.

Holmstrom, Bengt and Paul Milgrom. 1994. "The Firm as an Incentive System." *American Economic Review*. November, 84:4, pp. 972-991.

Isikoff, Michael and Mark Hosenball. 2004. "Bad Sourcing: U.S. Agencies May Have Relied on Fabricators and Saddam's Own Spies for Intelligence on Iraq." *Newsweek*. February 11, available at <<http://www.msnbc.msn.com/id/4244033/>>.

"Israel: What Went Wrong on October 6? The Partial Report of the Israeli Commission of Inquiry into the October War." 1974. *Journal of Palestine Studies*. Summer, 3:4, 189-207.

Jehl, Douglas. 2003. "Agency Belittles Information Given by Iraq Defectors." *New York Times* (final ed.). September 29.

Kreps, David M. 1990. "Corporate Culture and Economic Theory," in *Perspectives on Positive Political Economy*. James E. Alt and Kenneth A. Shepsle, eds. Cambridge: Cambridge University Press, pp. 90-143.

Lamont, Owen A. 2002. "Macroeconomic Forecasts and Microeconomic Forecasters." *Journal of Economic Behavior & Organization*. July, 48:3, pp. 265-280.

Lazear, Edward P. 1989. "Pay Equality and Industrial Politics." *Journal of Political Economy*. June, 97:3, pp. 561-580.

Lazear, Edward P. and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy*. October, 89:5, pp. 841-864.

Maskin, Eric, Yingyi Qian and Chenggang Xu. 2000. "Incentives, Information, and Organizational Form." *Review of Economic Studies*. April, 67:2, pp. 359-378.

Milgrom, Paul and John Roberts. 1988. "An Economic Approach to Influence Activities and Organizational Responses." *American Journal of Sociology* 94:Supplement, pp. S154-S179.

Milgrom, Paul and John Roberts. 1990. "Bargaining Costs, Influence Costs and the Organization of Economic Activity," in *Perspectives on Positive Political Economy*. James E. Alt and Kenneth A. Shepsle, eds. Cambridge: Cambridge University Press, pp. 57-89.

Mullainathan, Sendhil. 2000. "Thinking Through Categories." Mimeo, MIT.

Mullainathan, Sendhil and Andrei Shleifer. Forthcoming. "The Market for News." *American Economic Review*.

National Commission on Terrorist Attacks Upon the United States. 2004. *The 9/11 Commission Report: Final Report of the National Commission on Terrorist Attacks Upon the United States*. Washington, D.C.: Government Printing Office.

Ottaviani, Marco and Peter Sorensen (2000). "Her Behavior and Investment: Comment." *American Economic Review*, 90:30, pp. 695-704.

Posner, Richard A. 2005. *Preventing Surprise Attacks: Intelligence Reform in the Wake of 9/11*. Lanham: Rowman & Littlefield.

Prendergast, Canice. 1993. "A Theory of 'Yes Men.'" *American Economic Review*. September, 83:4, pp. 757-770.

Prendergast, Canice. 1999. "The Provision of Incentives in Firms." *Journal of Economic Literature*. March, 37:1, pp. 7-63.

Prendergast, Canice. 2002. "The Tenuous Trade-Off between Risk and Incentives." *Journal of Political Economy*, 110:5, pp. 1071-1102.

Prendergast, Canice, and Lars A. Stole. 1996. "Impetuous Youngsters and Jaded Old-timers: Acquiring a Reputation for Learning." *Journal of Political Economy*. December, 104:6, pp. 1105-1134.

Radner, Roy. 1992. "Hierarchy: The Economics of Managing." *Journal of Economic Literature*. September, 30:3, pp. 1382-1415.

Rajan, Raghuram and Julie Wulf. 2003. "The Flattening Firm: Evidence from Panel Data on the Changing Nature of Corporate Hierarchies." NBER Working Paper No. W9633, April.

Reiff, David. 2003. "Blueprint for a Mess." *New York Times* (final ed.). Magazine Desk, November 2.

Roberts, John. 2004. *The Modern Firm: Organizational Design for Performance and Growth*. Oxford: Oxford University Press.

Sah, Raaj Kumar and Joseph E. Stiglitz. 1986. "The Architecture of Economic Systems: Hierarchies and Polyarchies." *American Economic Review*. September, 76:4, pp. 716-727.

Scharfstein, David S. and Jeremy C. Stein. 1990. "Herd Behavior and Investment." *American Economic Review*. June, 80:3, pp. 465-479.

Skaperdas, Stergios. 1992. "Cooperation, Conflict, and Power in the Absence of Property Rights." *American Economic Review*. September, 82:4, pp. 720-739.

Stiglitz, Joseph E. 2001. Nobel Prize Lecture. "Information and the Change in the Paradigm of Economics." Available at <<http://nobelprize.org/economics/laureates/2001/stiglitz-lecture.pdf>>.

Welch, Ivo. 1992. "Sequential Sales, Learning, and Cascades." *Journal of Finance*. June, 47:2, pp. 695-732.

Williamson, Oliver E. 1967. "Hierarchical Control and Optimum Firm Size." *Journal of Political Economy*. April, 75:2, pp. 123-138.

Zwiebel, Jeffrey. 1995. "Corporate Conservatism and Relative Compensation." *Journal of Political Economy*. 103:1, pp. 1-25.