12

Process, Responsibility, and Myron's Law

Paul M. Romer

In the wake of the financial crisis, any rethinking of macroeconomics has to include an examination of the rules that govern the financial system. This examination needs to take a broad view that considers the ongoing dynamics of those rules. It will not be enough to come up with a new set of specific rules that seem to work for the moment. We need a system in which the specific rules in force at any point in time evolve to keep up with a rapidly changing world.

A diverse set of examples suggests that there are workable alternatives to the legalistic, process-oriented approach that characterizes the current financial regulatory system in the United States. These alternatives give individuals responsibility for making decisions and hold them accountable. In this sense, the choice is not really between legalistic and principle-based regulation. Instead, it is between process and responsibility.

The Dynamics of Rules

The driving force of economic life is the nonrivalry of ideas. *Nonrivalry* means that each idea has a value proportional to the number of people who use it. Nonrivalry creates a force that pushes for increases in the scale of interaction. We see this force in globalization, which relies on flows of goods to carry embedded ideas to ever more people. We see it in digital communication, which allows the direct sharing of ideas among ever more people. We see it in urbanization, which allows us to share ideas in face-to-face exchanges with ever more people.

A new slant on an old saying expresses the updated essence of nonrivalry of a technological idea: give someone a fish, and you feed them for a day; teach someone to fish, and you destroy another aquatic

ecosystem. This update reflects what has happened throughout most of human history and warns that we need more than new ideas about technology to achieve true progress.

We need to broaden our list of ideas to include the rules that govern how humans interact in social groups (rules like those that limit the total catch in a fishery). *Rules* in this sense means any regularities of human interaction, regardless of how they are established and enforced. Finding good rules is not a one-time event. As academics, policymakers, and students of the world, we need to think about the dynamics of both technologies and rules.

To achieve efficient outcomes, our rules need to evolve as new technologies arrive. They must also evolve in response to the increases in scale that nonrivalry induces. Finally and perhaps most important, they also need to evolve in response to the opportunistic actions of individuals who try to undermine them. Myron Scholes once captured this last effect in a statement he made in a seminar, a statement that deserves to be immortalized as Myron's law:

Asymptotically, any finite tax code collects zero revenue.

His point was that if there is a fixed set of rules in something like a tax code, clever opportunists will steadily undermine their effectiveness. They will do this, for example, by changing the names of familiar objects to shift them between different legal categories or by winning judicial rulings that narrow the applicability of the existing rules.

In sum, rules have to evolve in response to three distinct factors—new technologies, increases in the scale of social interaction, and opportunistic attempts at evasion. Any social group has higher-level rules—*metarules*—that determine how specific rules evolve. The metarules that govern the tax code, for example, allow for changes through legislation passed by Congress, regulations written by the Internal Revenue Service, and rulings handed down by courts. In some domains, the three forces that call for more rapid change in the rules may operate with greater force. In those domains, we presumably want to rely on different metarules.

Why Rules Lag Behind

As the number of people who use the Internet has increased, the rules that govern behavior have lagged far behind actual practice. This case

offers helpful illustrations about the general problem that we face ensuring that rules keep up.

New technologies are part of the problem. Digital communication has created many new possibilities for criminal activity that crosses national borders. Our systems of criminal investigation and prosecution, which are based on geographical notions of jurisdiction, are ill suited to this new world.

Scale also has an independent effect. Email is based on a set of rules that worked well when dozens of academics were communicating with each other. These informal rules were based on norms and reputation, so the Internet protocol and associated protocols for managing email failed to include even the most basic protections. Now that the Internet has scaled from dozens of people to billions, different rules are needed. For example, there is no built-in way for the recipient of an email to be sure about the identity of the sender. In a "spear-phishing" attack, an email is carefully tailored to resemble the authentic emails that the recipient normally receives. Because none of the usual warning signs are present (there are no offers of millions of dollars stranded in a stranger's bank account), the recipient is more likely to open an attachment with malicious code. Even RSA, a company whose business revolves around computer security, was compromised through this kind of attack.

Engineers at the Internet Engineering Task Force, a loosely defined voluntary organization with little formal authority, are the rule setters for the Internet. In 1992, they began to work on improving security protocols. They devised a patch called IPsec that reverse-engineered some basic security measures into the existing protocol. They also developed an update to the basic Internet protocol, known as IPv6, that has built-in support for IPsec. The basic specifications for these protocols were completed in 1998. Unfortunately, larger scale not only creates the need for better security but also makes it much harder to implement a change in the rules. The adoption of both sets of protocols has been held back by coordination problems among large numbers of users and vendors.

Even if these protocols are widely adopted, new attacks will still emerge. Bigger scale means that traditional mechanisms like reputation no longer operate and that more people are working to undermine and subvert all the existing security measures. Because a new vulnerability is a nonrival good that can be shared among predators, an increase in scale

can increase the rate at which predators circumvent any given security system.

Financial Markets

Rules in financial markets need to evolve for all of the reasons identified above. Technology is creating entirely new opportunities—for example, in high-frequency electronic trading systems. The scale of financial markets continues to grow, and private actors in these markets will surely seek clever ways to evade the intent of existing rules. The gains from opportunism in these markets are so large that the total amount of human effort directed at evading the rules will presumably be at least as large as that devoted to a low-return activity like cybercrime.

Electronic transactions were supposed to offer liquid markets and unified prices that can be accessed by everyone, but they have not lived up to this promise because they have also created new opportunities for manipulation. For example, some firms now submit and withdraw very large numbers of electronic quotes within milliseconds in a practice known as *quote stuffing*. It is not clear what the intent of these traders is, but it is clear is that any electronic trading system will have capacity constraints in computation and communication. Any system will therefore be subject to congestion. In the May 2010 stock market flash crash, congestion added to the anomalous behavior that firms were observing, and this apparently encouraged many high-frequency traders to stop trading, at least temporarily. This seems to have contributed to the temporary sharp fall in prices.

Quote stuffing could be one of many different strategies that traders use to influence local congestion and delays in the flow of information through the trading system. These, in turn, could affect liquidity as they did during the flash crash. As a result, transactions could take place at prices that depart substantially from those that prevailed just before or just after they occurred.

After an extensive analysis, the Securities and Exchange Commission (SEC) reported that quote stuffing was not the source of the cascade of transactions that overwhelmed the systems during the flash crash. The SEC is still equivocating about whether this particular practice is harmful

and, more generally, about systemic problems that high-frequency traders may be causing. Even if it had tried to address the specific practice of quote stuffing, the type of rule that had first been mooted—forcing traders to wait 50 milliseconds before withdrawing a quote that they had just submitted—would probably have been too narrow to limit the many other strategies that could be used to generate congestion or influence liquidity.

It seems implausible that the kind of behavior that occurred in the flash crash is an inevitable consequence of electronic trading. (But if it is, it seems implausible that the switch to electronic trading has brought net welfare benefits for the economy as a whole.) One year later, it also seems implausible that any of the changes implemented so far has fully addressed the underlying issue. Individual stocks continue to suffer from instances where trades take place at prices that are dramatically different from those that are prevailing seconds before or seconds later.

After the flash crash, trades were canceled if they took place at prices that differed from a reference price by more than a discretionary threshold, set in that particular case as a 60 percent deviation. Under new rules that try to be more explicit, transactions for some individual stocks will be allowed to stand if they take place at prices within 10 percent of the a reference price. In a multistock event, where many prices move together, the band of acceptability widens to 30 percent. Some have criticized these new rules because they still allow some discretion in setting the reference price. Others have expressed concern about the potential for manipulation that could intentionally trigger the looser rules that apply in a multistock event.

As the discussion below about rule making at Occupational Safety and Health Administration (OSHA) shows, even in a simple setting it is difficult to develop rules in a timely fashion that meet legal standards for clarity and do so following procedures that meet legal standards for due process. The Security and Exchange Commission's attempts to clarify the rules for breaking trades suggest that it is much harder to live up to these standards in a complicated and dynamic context. The SEC seems to have settled for a rule-setting process that leaves ample room for opportunism for extended periods of time. Perhaps some other, less legalistic approach deserves consideration.

Process versus Responsibility in Other Domains

One way to think about how the metarules that govern financial regulation might be adjusted so that the system can respond more quickly is to examine a broad range of social domains and observe the outcomes under alternative metarules. Here are four influential organizations in the United States that set rules and a specific goal that each organization's rules try to promote:

· Federal Aviation Administration (FAA): flight safety
· Federal Reserve: stable economic activity
· U.S. Army: combat readiness
· Occupational Safety and Health Administration (OSHA): worker safety

The Federal Aviation Administration works in a domain with the potential for rapid technological evolution. It has responsibility for passenger airplanes, which are among the most complex products ever developed. It approaches its task of ensuring flight safety with rules that specify required outcomes but that are not overly precise about the methods by which these outcomes are to be achieved. This is one way to interpret what principle-based regulation should look like. In practice, this means that some person must have responsibility for interpreting how any specific act, in a specific situation, either promotes or detracts from the goal that is implicit in the principle. That is, someone has to take responsibility for making a decision.

The general requirement that the FAA places on a new plane is that the manufacturer demonstrate to the satisfaction of its examiners that the new airplane is airworthy. The examiners use their judgment to decide what this means for a new type of plane. Within the FAA, the examiners are held responsible for their decisions. This changes the burden of proof from the regulators of a new technology to the advocates of the technology and gives FAA examiners a large measure of flexibility.

This approach stands in sharp contrast to one based on process. There is no codified process that a manufacturer can follow and be guaranteed that a new plane will be declared airworthy. Nor is there a codified process that the FAA examiners can follow in making a determination

about airworthiness. There is no way for them to hide behind a defense that they "checked all the boxes" in the required process.

One obvious requirement for a plane to be airworthy is that the airframe is sufficiently strong. There are no detailed regulations that specify the precise steps that a manufacturer must use to make a plane strong or show that it is strong. For example, there are no regulations about the size or composition of the rivets that hold the skin on the air frame, nor should there be. On an airplane like the Boeing 787, which is made of composite materials, there are no rivets. Instead, as part of the general process of establishing airworthiness, the employees of the FAA have technical expertise in areas like materials science and testing procedures and are responsible for making a judgment about how to test a particular design and determine whether it is sufficiently strong.

Moreover, because new information about an airframe can emerge for decades after it enters into service, the granting of a certificate of airworthiness is always provisional. Operators of aircraft are required to report evidence that emerges over time that might be relevant to airworthiness. At any time, the FAA can withdraw a plane's airworthiness certificate or mandate changes that must be made to an aircraft for it to continue to be airworthy. No judicial proceeding is required. There is no appeal process for an owner that unexpectedly receives an airworthiness directive that mandates an expensive modification. There is no way to get a judge to issue an injunction that would let the plane keep flying because the FAA has not satisfied some procedural requirement.

It is also clear that the rate of innovation in technologies is a choice variable along with the rate of innovation in the rules. If social returns are maximized when technologies and rules stay roughly in sync, good metarules might require that those who develop new technologies also have to develop the complementary rules before the new technologies can be implemented. A larger plane such as the Airbus 380 will generate more air turbulence in its wake. This means that the FAA has to implement new rules about the spacing between planes that follow each other on a flight path. The FAA will not let a plane like the Airbus 380 fly until the manufacturer has demonstrated the size of its wake and the FAA has had time to put in place new systemwide rules about separation. This is the polar opposite of the approach that the SEC takes with regard

to the introduction of major changes in the architecture of the electronic trading system.

The FAA implements a system based on individual responsibility by organizing itself as a hierarchy. People at a higher level can promote and sanction people at lower levels based on how well they do their jobs. At the top of the hierarchy, the secretary of transportation and the administrator of the FAA are appointed by the president and confirmed by the Congress, both of which are held accountable by the electorate.

The Federal Reserve, like every other central bank, is also organized as a hierarchy. Its leaders are held accountable by democratically elected officials who specify a mandate. In their day-to-day decisions, the employees at lower levels in the hierarchy have a lot of freedom to take actions that will achieve the organization's mandate. They are rewarded or punished based on the judgment of those one level higher in the hierarchy. There is little scope for the legislature to micromanage decisions, and there is no judicial review of the process by which decisions are made. As was seen in the financial crisis of 2009, this kind of system allowed for a much quicker response than the parallel mechanism involving legislation passed by Congress. The Fed's response to the failure of Long Term Capital Management also showed that it could manage what amounted to a bankruptcy reorganization far more quickly than a court could.

 Like the Fed and the FAA, the U.S. Army is run as a hierarchy with accountability at the top to elected officials. After a period during the 1970s when racial tensions in the army were seriously undermining its effectiveness, the leaders of the army decided that better race relations were essential for it to meet its basic goal of combat readiness. In less than two decades, they remade the organization. Writing in 1996, the sociologists Charles Moskos and John Butler observed that among large organizations in the United States, the army was "unmatched in its level of racial integration" and "unmatched in its broad record of black achievement" (2). To illustrate how different the army was from more familiar institutions such as the universities where they worked, Moskos and Butler (1996, 3) tell this story:

Consciousness of race in a nonracist organization is one of the defining qualities of Army life. The success of race relations and black achievement in the Army revolves around this paradox. A story several black soldiers told us at Fort Hood,

Texas, may help illustrate this point. It seems that one table in the dining facility had become, in an exception to the rule, monopolized by black soldiers. In time, a white sergeant came over and told the blacks to sit at other tables with whites. The black soldiers resented the sergeant's rebuke. When queried, the black soldiers were quite firm that a white soldier could have joined the table had one wished to. Why, the black soldiers wondered, should they have to take the initiative in integrating the dining tables?

The story has another remarkable point—that a white sergeant should take it on himself to approach a table of blacks with that kind of instruction. The white sergeant's intention, however naive or misdirected, was to end a situation of racial self-segregation. Suppose that a white professor asked black students at an all-black table in a college dining hall to sit at other tables with whites. This question shows the contrast between race relations on college campuses and in the army.

The system in the army makes such individuals as the sergeant in this story are held responsible for the state of race relations in any unit they supervise. This system holds them responsible for both their decisions and accomplishments, through occasional ad hoc review of their decisions by superior officers and through more formal decisions about promotion to a higher rank. Any particular decision like that of the sergeant in the story could easily be second-guessed, but the system as a whole has clearly been effective at achieving both integration and good race relations. Both direct judicial intervention in the operation of public school systems and the combination of legislation and regulations that guide behavior on university campuses have been far less successful.

The approaches to safety at the FAA, to macroeconomic stabilization at the Fed, and to race relations in the army all stand in sharp contrast to the legalistic, process-centered approach to safety followed by OSHA. To improve safety on construction sites, which have a bad safety record, OSHA follows a detailed process that leads to the publication of specific regulations such as these:

1926. 1052(c)(3)
The height of stair rails shall be as follows:
1926. 1052(c)(3)(i)
Stair rails installed after March 15, 1991, shall be not less than 36 inches (91.5 cm) from the upper surface of the stair rail system to the surface of the tread, in line with the face of the riser at the forward edge of the tread.
1926. 1052(c)(3)(ii)
Stairrails installed before March 15, 1991, shall be not less than 30 inches (76 cm) nor more than 34 inches (86 cm) from the upper surface of the stair rail

system to the surface of the tread, in line with the face of the riser at the forward edge of the tread.

These regulations are enforced by OSHA inspectors, who can issue citations that lead to fines and that can then be challenged in court. The regulations are supplemented by guidance about enforcement. For example, in the early 1990s, someone also added a note in the *Construction Standard Alleged Violations Elements (SAVE) Manual* that guided OSHA inspectors on how to apply these regulations on stair rails:

NOTE: Although 29 CFR 1926.1052(c)(3)(ii) sets height limits of 30″–34″ for stairways installed before March 15, 1991, no citation should be issued for such rails if they are 36″ maximum with reference to 29 CFR 1926.1052(c)(3)(i).

This change in enforcement patterns avoids the awkward situation in which a 35-inch-high rail could be cited either for being too low or for being too high depending on when it was installed, although it still leaves a puzzle about why a 38-inch-high rail might still be cited if it had been installed too early.

It is tempting to ridicule regulations like these, but it is more informative to adopt the default assumption that the people who wrote them are as smart and dedicated as the people who work at the FAA. From this, it follows that differences in what the two types of government employees actually do must be traced back to structural differences in the metarules that specify how their rules are established and enforced. The employees at the FAA have responsibility for flight safety. They do not have to adhere to our usual notions of legalistic process and are not subject to judicial review. In contrast, employees of OSHA have to follow a precise process specified by law to establish or enforce a regulation. The judicial checks built into the process mean that employees at OSHA do not have any real responsibility for worker safety. All they can do is follow the process.

One possible interpretation of the regulations about stair rails is that the regulations once specified a maximum height of 34 inches and that new evidence emerged showing that a higher rail would be safer. As they considered new rules they could propose, the regulation writers faced the question of what rules to apply to stairways that had been installed in the past. Rather than make an ex post change to the regulations for existing stairs, they may have chosen instead to stick to the principle that the regulations that were in force when a stairway was installed would

continue to apply to that stairway but to suspend enforcement for some violations.

The caution about ex post changes in the regulations may derive in part from a concern about judicial review of the new rules. Or it could have come from a concern about judicial review of penalties that had already been assessed or violations under the old rules that would no longer be violations under the new rules. The change in enforcement at least made sure that no judge saw cases where 35-inch-high rails were sometimes cited for being too high and sometimes for being too low.

You can get some sense of how difficult it is to be precise in writing rules by digging into an area like this. From published inquiries that OSHA received, it seems that the decisions here were complicated by ambiguity about the rules for handrails, which a person uses as a grip and should therefore not be too high, and stair rails, which mark the top of a barrier designed to prevent falls and which therefore should not be too low. The top of a stair rail might be but need not be a handrail. It looks as though the rules morphed over time to distinguish more explicitly between the two types of rails.

It is striking that safety officers for construction firms who wrote to OSHA for clarifications about apparent discrepancies between different sections of the regulations waited four to six months to receive answers. (One wonders what happened at the construction site during those many months.)

Even more striking is the fact that the rules cited here were first proposed in 1990 or 1991, but judging from a 2005 notice in the *Federal Register* calling once again for comments, they did not come into force until sometime after 2005. (The notice in 2005 makes a brief reference to other agency priorities that took precedence over the rules for stair rails.) This required the application of a further enforcement instruction that a stair rail that conforms to the proposed regulation for stairs built after 1991 but that violates the existing regulations (which were not changed for another fifteen years) would be treated as a de minimus violation and would not result in an enforcement action.

The principle-based approach to the regulation of air safety lacks all of the procedural and legal protections afforded by the process of OSHA, but in terms of the desired outcomes, the FAA approach seems to work better. Air travel is much safer than working on a construction site. The

Fed and the army also seem to have been much more effective in addressing complicated challenges. Despite the more extensive judicial protections afforded the construction firms under the OSHA process, firms find the process infuriating. Construction sites are still very dangerous places to work.

Conclusion

People from the United States take pride in a shared belief that theirs is "a nation of laws, not of individual men and women." Taken literally, this claim is nonsense. Any process that decides what kind of planes can carry passengers, what to do during a financial panic, how people of different races interact, or how a construction site is organized will have to rely on decisions by men and women.

Because of combinatorial explosion, the world presents us with a nearly infinite set of possible circumstances. No language with a finite vocabulary can categorize all these different circumstances. No process that writes rules in such a language can cover all these circumstances. Laws and regulations always require interpretation. Giving judges a role to play in making these interpretations or reviewing them does not take people out of the process.

We could have a system in which individual financial regulators have the same kind of responsibility and authority as the sergeant in the cafeteria. If they saw behavior that looked harmful to the system, they could unilaterally stop it. We could have a system like the one we use to certify passenger aircraft, in which the burden of demonstrating that an innovation does not threaten the safety of the entire trading system rests on those who propose the innovation. In such a system, the people that the innovators would have to persuade could be specialists who would have the same kind of responsibility and authority as FAA examiners. The opportunists in the financial sector would presumably prefer to stay with an approach that emphasizes process, but this leaves the other participants in the sector at a relative disadvantage. More seriously, it leaves those outside the sector unprotected, with no one who takes responsibility for limiting the harms that the sector can cause.

The right question to ask is not whether people are involved in enforcing a system of rules but rather which people are involved and which

incentives they operate under. There may be some contexts where a legalistic approach like that followed by OHSA and the SEC has advantages, but we need to recognize that this approach is not the only alternative and that it has obvious disadvantages.

A careful weighing of the costs and benefits will involve many factors, but the factor that seems particularly important for the financial sector concerns time constants. As the OSHA example suggests, the legalistic process is inherently much slower than a process that gives individuals more responsibility. Moreover, clever opportunists can dramatically increase the delays and turn the legalistic approach into what Phillip K. Howard (2010) calls a "perpetual process machine."

Under this approach, rules for the financial sector will never keep up. The technology is evolving too quickly. The scale of the markets is enormous and continues to grow. There may be no other setting in which opportunism can be so lucrative. It is hard to understand why technologically sophisticated people devote any effort to committing cybercrime when the payoffs from opportunism in financial markets seem to be so much larger. If we persist with the assumption that a legalistic rule-setting process is the only conceivable one we could use to regulate financial markets, then the opportunists will thrive. We will settle into a fatalistic acceptance of systemic financial crises, flash crashes, and ever more exotic forms of opportunism.

"No one can predict how complicated software systems will behave" (except in airplanes). "You can't change behavior" (except in the army). "Financial systems are just too complicated to regulate" (except in countries like Canada, where instead of running a process, regulators take responsibility).

References

Howard, Philip K. 2010. "To Cut the Deficit, Get Rid of Surplus Laws." *Washington Post*, December 10.

Moskos, Charles C., and John Sibley Butler. 1996. *All That We Can Be: Black Leadership and Racial Integration the Army Way*. New York: Basic Books.