

Regulating AI: The Race Between Policymakers and the Rise of a New Interest Group*

Steven Callander[†]

April 27, 2025

Introduction

It is not unusual for the textbook version of an institution to differ from how it works in practice. This gap is nowhere greater than with politics. In the textbook version, policymaking reflects a careful balancing of benefits and harms in the public interest. In practice, politics is a brutal battle between interest groups, where fairness and the public interest play at best an incidental role.

The emergence of new technologies — and their possible regulation — bring this distinction to a sharp point. As a technology emerges, firms enter the market, and an industry develops around it. As the industry grows in the market, it also grows in stature in politics. The industry becomes an *interest group*. In the idealized version of politics, the industry would inject its expertise into policy deliberations and improve social welfare. In practice, the industry uses its political power to shape regulations in a way that is favorable to the industry itself.

For an established and stable technology, the solution to this problem is easy: impose socially-optimal regulation before the industry grows and obtains political strength. For new technologies like AI, the problem is that we don't yet know what those regulations should be. The policymaker can wait to learn, but by the time she identifies the optimal regulation, the AI industry may have grown so strong politically and it will be too late to stop it.

Regulating a new industry is a race. A race between policymakers and the rise of a new interest group. Policymakers race to learn about the technology and how to regulate it before the industry is able build political power and wield it shape regulation in its favor. In this chapter we describe the nature of this race and speculate on how it will develop. We ask: How can a socially-minded policymaker structure and compete in this race such that the winner

*I thank the audience and organizers of the NBER Economics of AI conference for stimulating discussion and an anonymous referee for helpful comments.

[†]Stanford University, Stanford Graduate School of Business; sjc@stanford.edu.

is society? This is a hard problem, to be sure, although not intractable. We sketch several paths that begin toward a solution.

Disruptive Technologies: Learning about Benefits and Harms

The impact of a new technology can often only be learned by doing. Only by developing and applying a technology does its power reveal itself and the harm it causes become evident. A challenge for policymakers is to learn about harms before they occur. If learning can occur in a lab, independent of industry growth, then a precautionary principle applies and regulators should wait to learn about harms before allowing the industry to flourish (Acemoglu and Lensman, 2024). In practice, lab learning is limited and most learning, particularly the impact a technology has on society, can only be learned by applying the technology in practice. Gans (2024) characterizes when it is optimal to let the industry grow despite the risks and, indeed, when it is optimal to speed up industry development to bring forward learning. Callander and Li (2024) introduce the possibility that the firms developing AI hold superior information about potential harms. They show that it is difficult for the policymaker to extract this information given the firms hope to profit from AI, and that this difficulty increases the more competition there is in the market.

The problem we address in this chapter is related yet different. In the papers cited above, the question is whether the policymaker can learn about harms before they occur. If she can, she can update regulations and avert the harms. The question of this chapter is whether the policymaker will still hold enough political power to update regulations by the time she learns whether harm will occur.¹ We follow Gans (2024) in supposing that learning is by doing and that the rate of learning depends on industry growth. To this we add the observation that industry growth also translates into growth in political power, and that the industry can use this political power to stop the policymaker from intervening in its use of AI.²

Political Outcomes are the Result of Interest Group Competition

An axiom of politics is that elected officials care about votes. Even public-minded policymakers must get reelected, and not all policymakers are that. The desire for votes need not create bad incentives. After all, satisfying citizens and winning their vote is the intent of representative democracy. The problem is that votes are not represented equally. If a voter is unaware of an issue, or if the issue will not swing her vote and she won't otherwise get

¹Gans (2024) studies two cases: when AI adoption is and is not reversible. In this chapter, the policymaker's ability to reverse AI adoption is endogenous, declining in time as the industry grows.

²The crypto industry provides an alarming example of the problem. The crypto industry leveraged the enormous market valuations of its largest firms to become the largest political contributor in the 2024 electoral cycle in the U.S.

involved, policymakers are free to ignore the voter’s preferences. Rather, it is those who care enough about an issue for their vote to hinge on how the policymaker acts who receive representation. In fact, it is not even necessary to be a voter to be represented. If an individual or group can deliver votes in some other way — via campaign donations, a positive message, or organization — then policymakers will represent their preferences. These are the *interest groups*. Olson (1971) taught us that the political process should be understood through the lens of interest groups. To predict a political outcome, Olson argued, we shouldn’t look too much at polls of public opinion, instead we should look at the weight of the interest groups on either side of the issue.

The political power of an interest group depends on several characteristics. Importantly, these characteristics are correlated with but are not in direct proportion to a group’s market or social power. To be powerful, an interest group must be organized. It must be able to act on its preferences and deliver votes to a policymaker. The number of votes that can be delivered, whether directly or indirectly, depends on (i) numbers — how many votes lie within the firm or industry and in its value chain? (ii) message — how sympathetic is the public to the interest group, (iii) money, and (iv) coverage — where are the votes located given the electoral rule in place?

As the AI industry grows, it will acquire political power to the extent that its market power translates into these characteristics. As it does so, a socially-minded policymaker’s ability to control policy will endogenously erode.

Politics Combined with Market Competition Generates Path Dependence in Technology

Politics not only creates a contest between policymakers and the AI industry, it creates competitive pressures within the industry. The combination of the two forces can generate a path dependence in the technology itself. This path dependence can lock in the technological approach adopted by early entrants, cutting off experimentation and innovation prematurely.

Path dependence emerges because first movers in the market are also first movers in politics. Later entrants must play catch up in the market as well as in politics. The need to catch up in politics can shape their market strategy and, indeed, their approach to the technology.

To see why, observe that the first movers use the political power they acquire to not only keep a policymaker off their back, but to keep at bay the threat of market competitors. If later entrants are poised to overtake them with new technological approaches, first movers can stop them by lobbying for standards that favor their own technological approach. This may not be good for AI broadly, but it is good for the first movers.

The strategic pressure on later entrants, therefore, is to hew closely to the technological

and market approach of the early entrants. To beat them on their own terms in the market so that political power cannot be turned against them. This pressure to conform leads to a technological path dependence that locks in whatever approach happened to be tried first, leaving other approaches unexplored.

Politics provides a novel source of technological lock-in. It complements existing explanations of path dependence and technological lock-in, such as learning-by-doing and the idiosyncrasy of early decisions (Bryan, 2017). Unlike lock-in from learning-by-doing, there are no redeeming features of lock-in from politics. When driven by politics, technological lock-in represents a pure inefficiency due to rent seeking that distorts and restrains the innovation process.

What Can be Done?

The problem we address here is not easy. In the face of extreme uncertainty, non-specialist policymakers must navigate a fine line between fostering innovation while not losing their authority over it. We sketch two possible paths to solving the problem. Both paths derive from the problem itself.

The first path is to drive a firmer wedge between markets and politics. If the industry cannot translate its market power into political power, then it can be allowed to flourish without fear that it will overwhelm the political process. There are several ways that the political process can be shielded from market power. One is to delegate AI regulation to an independent government agency. Although the bureaucracy always sits under legislative oversight, an agency is protected from political pressure by super-majority provisions in the legislature (e.g., the filibuster in the U.S. Senate) or other statutory requirements. This approach raises the possibility of capture of the regulator by the industry, so it is important which agency is granted authority and the nature of that authority.

One way to reduce the probability of regulatory capture is for the legislature to create a new agency dedicated to the regulation of AI. Imagine how the finance industry would look different today if the Consumer Finance Protection Bureau (CFPB) had been created before rather than after industry malpractice in the 2000's? Another, more minimalist, approach is to preemptively assign authority over AI to an existing agency. The overlapping jurisdiction of the multitude of government agencies, particularly at the federal level in the US, allows firms as well as activist groups to "venue shop" for a favorable agency. The legislature can retain control by clearly and preemptively delegating authority over AI. Do policymakers prefer that the agency in control of AI regulation be staffed by lawyers or by economists or even by social theorists? Or should those with technical training in the technology itself control regulation of the technology? This choice matters and it is one policymakers in the legislature have *ex ante* control over.

A second path to solving the problem is to target market power before it can be translated into political power. A blanket approach to limiting market power would be destructive to AI innovation. Fortunately, that is not necessary to solve the problem. It is sufficient to limit only the structure of market power so that the translation into political power is muted. As detailed above, the mapping from market to political power is not one-to-one. The AI industry can be allowed to flourish, with some firms even allowed to obtain substantial market power, so long as the form of that market power does not hold the elements of substantial political power. Regulating with this connection in mind means that regulation can be minimal and focused on shaping industry growth rather than limiting it.

This approach may involve regulation that separates the various functions of the industry into different firms so that the ingredients for political power don't accumulate. For example, the parts of the industry with the most money and resources could be separated from the parts of the industry with many jobs and a message that resonates with the public. This should be a consideration in the decision where to apply AI regulation — at the LLM or the app level — as it will shape how divisions within the industry develop.

In the best case, these interventions will shape the industry so that vibrant competition emerges within the AI industry, and that this translates into competition in the political domain as well. If the industry can be fractured politically, policymakers may obtain the best of both worlds: An industry without the unified strength to overwhelm the political system, and incentives that induce AI firms to reveal their private information about the technology through lobbying that will lead to more effective policymaking.

Whether policymakers can pull this off and solve their problem remains to be seen. This chapter points the way to a solution, one that allows the industry to develop and flourish while mitigating the political threat. The longer policymakers delay, the stronger the AI industry becomes, and the harder the problem becomes. The clock is ticking.

References

Acemoglu, D. and T. Lensman (2024). Regulating transformative technologies. *American Economic Review: Insights* 6(3), 359–376.

Bryan, K. A. (2017). The perils of path dependence. In J. Gans and S. Kaplan (Eds.), *Survive and Thrive: Winning Against Strategic Threats to Your Business*. Dog Ear Publishing.

Callander, S. and H. Li (2024). Regulating an innovative industry. Working paper.

Gans, J. S. (2024). How learning about harms impacts the optimal rate of artificial intelligence adoption. *Economic Policy*, 3–21.

Olson, M. (1971). *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press.