The Journal of Law, Economics, and Organization, 2023, **00**, 1–28
https://doi.org/10.1093/jleo/ewac025
Article

OXFORD

Approval regulation and learning, with application to timing of merger control

Marco Ottaviani[1,*], Abraham L. Wickelgren[2]

[1]Department of Economics, IGIER, and BISDA, Bocconi University, Milan, Italy
[2]School of Law, University of Texas at Austin, Austin, TX, USA

*Corresponding author: E-mail: marco.ottaviani@unibocconi.it

ABSTRACT

This article analyzes the optimal combination of ex ante and ex post regulation of an activity in a two-period model. Additional information about the sign and extent of the externality associated with the activity becomes available only once a private party undertakes the activity, but undoing the activity at that stage is costly. We characterize when the regulator should commit not to reevaluate the activity ex post. The case for ex post regulation is strengthened if the private party can signal its private information about the consequences of the activity, but it is weakened if the cost of undoing the activity can be manipulated.

1. INTRODUCTION

In 2014, both the Federal Trade Commission (FTC) and the European Commission (EC) approved Facebook's acquisition of WhatsApp for $19 billion, a landmark deal in digital markets and the largest one for Facebook. Six years later, the FTC is suing Facebook for anti-competitive behavior, citing also the 2012 takeover of Instagram. The FTC claims that Facebook was pursuing a strategy of acquiring apps in order to eliminate current and future competitive threats—as CEO Mark Zuckerberg wrote in a 2008 email, "it is better to buy than compete."

The stark change in direction by the FTC is one of many examples of a general backlash against Big Tech, as concerns around issues such as privacy, data sharing, and supranormal profits have been mounting in recent years. A simultaneous outage of Facebook, Instagram, and WhatsApp in October 2021 affected billions of people and is a clear reminder of the platforms' dominant role in society. The FTC's willingness to renege on its previous approval is a sharp example of how ex post regulation can be used to correct the errors committed in the ex ante stage. In addition, there have been calls to revisit prior clearance of Amazon's acquisitions of Whole Foods and Zappos as well as Google's acquisitions of DoubleClick and Nest (Patel 2020).

A key challenge faced by any regulator in ex ante approval regulation (of which merger review is one example) is the level of uncertainty about the effects of the activity prior to its taking place. The EC's decision in the WhatsApp case consisted of an analysis of hypothetical scenarios such as future data sharing between the two parties, WhatsApp potentially entering the online advertising market, and the integration between the two platforms (European Commission 2014).[1] As is typical for ex ante regulation, only by approving the merger could the agency confirm whether these activities will be completed and obtain information about their actual impact, in both magnitude and direction, on welfare.

In this article, we investigate a regulator's tradeoff between approving an activity ex ante based on little information and waiting to acquire more information. A similar tradeoff is at the core of competition policy, many forms of privacy and environmental regulations, as well as other types of approval regulation. We then analyze three additional issues raised by this tradeoff, each with clear applicability to the Facebook–WhatsApp merger and to other cases in digital markets. First, we show that ex post regulation can result in a commitment problem, in that the possibility of ex post regulation can lead firms to decide not to propose activities that, in expectation, are both profitable and increases social welfare. Second, ex post regulation can lead to a socially desirable discipline effect because firms deviate from profit-maximizing behavior to increase social welfare in order to avoid ex post reversal of profitable activities by the regulator. For example, we can justify Facebook's decision to reduce misinformation on WhatsApp by limiting excessive forwarding of messages as an attempt to gain goodwill with policymakers. Third, the threat of ex post regulation can lead firms to increase deliberately the social cost of undoing the activity.

The article highlights the key role of uncertainty in approval regulation. The level of uncertainty has increased in recent years given the fast pace of innovation in digital markets, where it is difficult for regulators to reliably predict future competition or privacy impacts, the evolution of technologies, and the actions of the players involved. As noted by Andrea Coscelli, Chief Executive of the UK Competition and Markets Authority (CMA), in 2019:

> Over the last decade, Amazon, Apple, Facebook, Google and Microsoft combined... have made over 400 acquisitions globally.... However only a handful of these mergers have been scrutinized by competition authorities, and none have been blocked.

It is possible that the greater uncertainty in digital markets compared to other industries has skewed the tradeoff towards ex post regulation, as predicted by the model. This increases the risk of having approved mergers that should have been blocked (Type II error). Argentesi et al. (2021) note that new entrants are relatively more valuable in digital markets because of network effects, the value of data, and the high barriers to entry characterizing these markets. Errors in approval are therefore becoming more costly. To the extent that learning about welfare effects in digital markets not only provided information about past activities but also reduced the uncertainty associated with current and future activities, it might be sensible to shift toward more ex ante regulation in the future.

[1] In the ruling, the EC concluded that "technical integration between WA and FB was unlikely to be as straightforward from a technical perspective as presented by third parties." Instead, the post-merger years saw significant integration between Facebook, WhatsApp, and Instagram (owned by Facebook and a leading company in the social network industry). In addition, the EC's conclusion was based on the premise that the different customer identifiers used by Facebook and WhatsApp were excessively difficult to match. This proved untrue, and the EC fined Facebook for providing inaccurate information. Nonetheless, the EC declared that they had considered this case too. (https://ec.europa.eu/commission/presscorner/detail/e/IP_17_1369)

Our analysis reveals that ex post control may have an undesirable chilling effect on the ex ante incentives of business to undertake socially desirable activities. In some cases, an activity would not be undertaken unless the government is able to commit not to undo it when it turns out to be socially detrimental. We identify situations in which the government benefits from committing to a regime based exclusively on ex ante review. In some instances, the government might be able to commit ex ante to an ex post standard that is not ex post optimal—we characterize situations in which it is optimal for the government to adopt a relatively lenient ex post standard to alleviate the chilling effect.

We extend the analysis to cases in which the new information regarding an activity's effects is observed only privately by the firm. The firm might then have incentives to distort its actions to influence the government's ex post regulatory decision. Our analysis uncovers instances in which this signaling distortion actually increases efficiency without loss of information.

In some instances, the firm might have an incentive to increase the cost of reversing the activity in order to reduce the occurrence of ex post control. For example, as the outage in October 2021 suggests, Facebook has begun integrating its platform with WhatsApp's, thus making a potential breakup costlier. To analyze this manipulation effect, we also extend the model to allow for endogenous reversibility costs.

Beyond the application to competition policy, our analysis is relevant for the wide set of regulatory approval processes put in place by governments to protect consumers from potential harm from private economic activities. There is often substantial uncertainty about the magnitude—and even the sign—of the externality generated by the private activity, but more precise information becomes available only once the activity is undertaken. At that stage, some of the damage has already been done and the activity becomes costly to reverse.

Our model broadly applies to the design of sequential approval processes for consumer protection in the health, safety, and environmental area. For example, consider an energy company interested in making an investment in an environmentally sensitive area (e.g., drilling for oil in the Alaskan wilderness) amid uncertainty about both the profitability of the investment and the ultimate environmental impact. Local government authorities will have much more precise information about how best to resolve this tradeoff once the activity has been undertaken for some time, and its actual effects have been observed. However, shutting down the project at that stage comes at a cost.

Next, consider an agricultural biotechnology corporation seeking approval for a genetically modified organism. Regulators such as the Environmental Protection Agency will know much more about the safety of the organism after it has been used. At that point, however, substantial sunk costs will have already been incurred in the manufacture and distribution of the organism. Similarly, consider a pharmaceutical firm introducing a new drug. While both the Food and Drug Administration (FDA) and the firm have some information from clinical trials, there remains a great deal of uncertainty about the longer-term safety of the drug. This uncertainty will only be resolved if the FDA approves the drug and many have experienced its effects. The tradeoff between ex ante and ex post approval regulation analyzed here is relevant also for professional licensing and urban planning permits.

Lastly, consider a bank offering a complex financial product, which might be optimal for some sophisticated consumers, but might harm other unsophisticated consumers who are misled into purchasing this product. A regulator, such as the Consumer Financial Protection Bureau, might want to approve the product if the number of misled consumers is small, but it will have much better information about the impact after the product has been sold for a while.

The article proceeds as follows. Section 2 formulates our baseline model with symmetric learning about the consequence of the activity. Section 3 analyses the main tradeoff between flexibility and commitment. Section 4 extends the model to allow the firm to observe privately the consequences of the activity. Section 5 analyzes the firm's incentives to increase the cost of reversing the activity so as to reduce the occurrence of ex post control. Section 6 discusses how allowing the regulator to approve the activity subject to some conditions might affect the analysis. Section 7 develops applications to merger review and environmental regulation. Section 8 places our contribution within the literature. Finally, Section 9 discusses our main findings and concludes. All proofs not presented in the text are collected in the Appendix.

2. BASELINE MODEL WITH SYMMETRIC LEARNING

A firm, F, and a regulator, R, play the following game. In period 0, F is considering undertaking some activity, A, that generates a known profit π in period 1 and $n\pi$ in period 2, where n represents the discounted duration of period 2 relative to period 1. The activity also generates an externality. In period 0, the magnitude of the externality is not known. As a result, the total social welfare effect, which we call θ, is uncertain in period 0. We will refer to θ as the state. The state is distributed according to distribution function $G(\theta)$ with associated support $[\underline{\theta}, \overline{\theta}]$ and density function $g(\theta)$.[2] Note that $\theta - \pi$ measures the externality associated with activity A. We will confine attention to cases in which profit is always weakly monotonic in the state. That is, $\pi'(\theta) \leq 0$ for all θ or $\pi'(\theta) \geq 0$ for all θ. We assume $E[\pi] = \int_{\underline{\theta}}^{\overline{\theta}} \pi(\theta)g(\theta)d\theta > 0$, so the activity is profitable in expectation for the firm.

For example, if the activity is a merger, there may be uncertainty about the degree to which the merger increases market power, in which case larger θ represents a smaller increase in market power, making $\pi'(\theta) \leq 0$. Alternatively, when the efficiencies from the merger are uncertain, we would expect $\pi'(\theta) \geq 0$ because greater efficiencies increase both welfare and profits. In the case of environmental regulation affecting energy production, greater existing energy supplies (without the activity) suggest that both welfare and profit from the activity are smaller, so $\pi'(\theta) \geq 0$. When applying the model to regulation of consumer products such as balloon mortgages that may benefit some consumers while harming naive consumers, we could specify that greater consumer sophistication increases welfare but lowers profits, so $\pi'(\theta) \leq 0$. This monotonicity assumption reflects a significant number of interesting cases, but does not capture every regulatory situation.

We assume that the activity is subject to regulatory approval, but not liability. That is, the firm F can only undertake activity A if it is approved by a regulator R (who maximizes social welfare). The state measures R's payoff, equal to social welfare, from activity A, which is given by θ in period 1 and $n\theta$ in period 2. In each period in which the activity is not in place, we normalize the firm's payoff and social welfare to zero. To avoid trivial cases, we focus on situations in which there is genuine uncertainty about the social desirability of activity A by assuming that $0 < G(0) < 1$, so that $\underline{\theta} < 0 < \overline{\theta}$.

In the baseline version of the model, we assume that the value of θ becomes publicly known in period 1 only if the firm undertakes A in period 0, so that the state remains unknown if the activity is not undertaken. The regulator has the choice of blocking A in period 0 or waiting until period 1 to learn θ (assuming F undertook A in period 0) and mandate F to undo A. Once A has been undertaken in period 0, F must bear a cost equal to $k \geq 0$

[2] Equivalently, G can be reinterpreted as the posterior distribution generated from a noisy signal about the effects of the action.

to undo A in period 1. This cost also enters the payoff of R because the firm's payoff is part of social welfare. Parameter k measures the extent of irreversibility associated with the activity.

To summarize, the timing of the game is as follows. At period 0:

- R decides whether to allow or block activity A.
- If R allows activity A, F decides whether to undertake activity A.
- If activity A is not undertaken by F, the game is over and both R and F obtain the baseline payoff of 0.

If activity A is undertaken by F at period 0, at period 1:

- θ becomes publicly known.
- R obtains payoff θ, and F obtains payoff π.
- R decides whether F should undo activity A, in which case F bears cost k and the game is over.

If activity A is in place at the end of period 1, at period 2:

- R obtains payoff $n\theta$ and F obtains payoff $n\pi$.

3. FLEXIBILITY VERSUS COMMITMENT: CHILLING EFFECT

We begin by considering the ex post optimal decision made by the regulator in period 1. If the firm undertakes activity A in period 0, in period 1, the regulator only undoes the activity going forward if the future social cost of the activity, $-n\theta$, exceeds the social cost of undoing the activity, k. If instead $\theta \geq -k/n$, R allows the activity to continue. The regulator's ex post decision reflects the information that becomes available in period 1, while also taking into account that reversing the activity is costly. When θ lies in the interval $(-k/n, 0)$, in hindsight the regulator should have blocked the activity in period 0; see Figure 1. We assume that π is weakly monotonic in θ, although we allow for it to be both non-decreasing or non-increasing.

For simplicity, we make the following assumption so that the firm does not find it optimal to voluntarily undo the activity in period 1 after learning θ.

Condition 1. $\pi(-k/n) \geq -k/n$ and $\pi(\overline{\theta}) \geq -k/n$

If $\pi(\theta)$ is monotonically nondecreasing, the first part of the condition guarantees that the firm will continue the activity whenever the regulator approves it ex post. If instead $\pi(\theta)$ is monotonically nonincreasing, the second part of the condition guarantees continuation by the firm since the firm's lowest profit state is at $\overline{\theta}$, where the regulator will certainly allow the action to continue. Note that the second part of the condition implies that there exists a $\check{k} = -n\pi(\overline{\theta})$ that provides a lower bound on the admissible k in our model when $\pi(\theta)$ is monotonically nonincreasing. If k is too small relative to the firm's minimum profit when profit is negatively correlated with social welfare, then the firm would voluntarily undo the activity in the states in which the regulator most wants the activity to continue.[3] If $\pi(\theta)$ is monotonically nondecreasing, then the lower bound is k^* satisfying $\pi(-k^*/n) = -k^*/n$.

[3] If these conditions are not satisfied, allowing the firm to undo A voluntarily in period 1 introduces a second option value to the model. The analysis of the model in the presence of this second option value is of independent interest but would distract from the main message of this article.

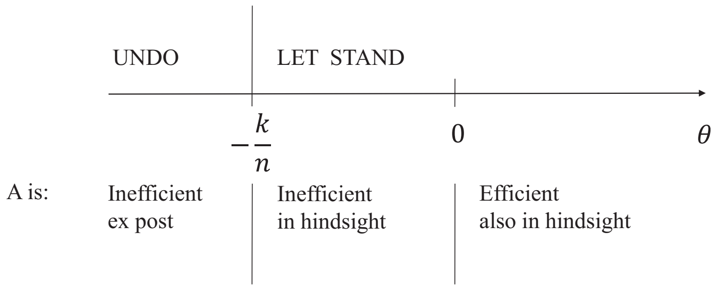

Figure 1. Ex post control.

The expected ex ante social welfare is then

$$W[k, n, g(\theta)] := E[\theta] + n[1 - G(-k/n)]E[\theta|\theta \geq -k/n] - kG(-k/n) \tag{1}$$

where $E[\theta] = \int_{\underline{\theta}}^{\overline{\theta}} \theta g(\theta) d\theta$ is the unconditional expected social welfare from activity A and $-k/n$ is the *ex post threshold* of the state below which the regulator undoes A in period 1.

To analyze R's decision in period 0 whether to block A, we begin by supposing that the differential profits, π, are sufficiently large that F undertakes A regardless of R's policy. Our first result describes the regulator's ex ante optimal decision.

Proposition 1. If it is profitable for the firm to undertake the activity in period 0, it is optimal for the regulator to block the activity ex ante if and only if $E[\theta] < 0$ and the cost of undoing the activity in period 1 is greater than a certain threshold, $\hat{k} \geq 0$.

Proof. It is optimal for R to block A in period 0 if and only if $W[k, n, g(\theta)] < 0$. As illustrated in Figure 2, W is decreasing in k, with

$$\frac{\partial W}{\partial k} = -G\left(-\frac{k}{n}\right) < 0 \tag{2}$$

for $k < -n\underline{\theta}$; W reaches $(1 + n)E[\theta]$ at $k = -n\underline{\theta}$; and it remains constant for $k > -n\underline{\theta}$. Therefore, W reaches its maximum at the lower bound $k = 0$. If $W[0, n, g(\theta)] \geq 0$ and $E[\theta] < 0$, then there exists $\hat{k}$ such that $W[\hat{k}, n, g(\theta)] = 0$. In this case, in period 0 the regulator blocks A whenever $k \geq \hat{k}$, and does not block otherwise. If $E[\theta] > 0$, then $W > 0$ for all k, and it is never optimal to block A in period 0. If $W[0, n, g(\theta)] < 0$, it is optimal to block A in period 0 for all k. *Q.E.D.*

Clearly, it is never optimal to block an activity ex ante that increases expected welfare. When instead the expected social welfare effect of the activity is negative, this proposition says the activity should be blocked ex ante if and only if the cost of undoing it later, k, is large enough. When the cost k is low, it is optimal to allow the activity in period 0 even though the activity results in an expected reduction in social welfare. In this case, by allowing the activity to go ahead, R obtains more information and is then able to decide in period 1 whether to undo the activity. Intuitively, the smaller is the cost of undoing the activity when it turns out to reduce social welfare, the larger is the option value of learning the precise welfare effect of the activity.

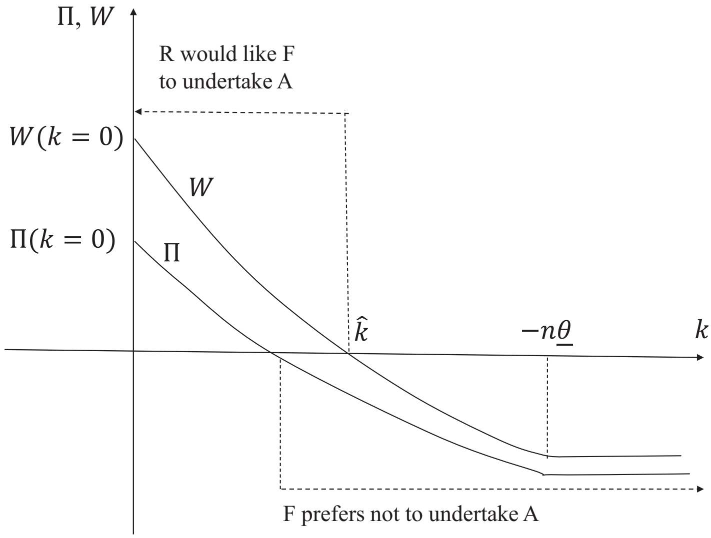

Figure 2. Value of commitment.

Next, we show that the parameter region for which it is optimal for R to block A ex ante shrinks as the weight of the second period payoff increases. This follows because the option value of more accurate information increases when the weight of the period in which one can use that information increases. Similarly, the option value of ex post review increases with a first-order stochastic shift and with a mean-preserving spread in the distribution of the state. If the state is generally more favorable, then there is less risk and more benefit from allowing the activity to learn the actual state. In addition, if there is greater variance in the possible states, then the value of learning the actual state is larger.

Proposition 2. An increase in the duration, n, of the second period, a first-order stochastic increase in the distribution of θ, and a mean-preserving spread in the distribution of θ lead to an increase in the ex ante control threshold for the reversal cost, $\hat{k}$.

3.1 Ex post control

So far we supposed that the firm is willing to undertake activity A even when it is subject to ex post review by the regulator. Next, we illustrate cases in which the threat of ex post review makes it unprofitable for the firm to undertake the activity. In particular, we characterize a range for the reversal cost for which R's ability to undo A ex post (when A turns out to be socially undesirable) deters the firm from undertaking A in the first place, even when A is ex ante socially beneficial.

The firm's expected profit from activity A in period 0, given the regulator's ex post approval policy, is

$$\Pi[k, n, g(\theta)] := E[\pi] + n[1 - G(-k/n)]E[\pi|\theta \geq -k/n] - kG(-k/n). \tag{3}$$

When this is negative, the firm is deterred from undertaking activity A. If so, R might benefit from committing *not* to review the activity ex post. This happens if both the expected welfare and expected profit (given commitment not to reverse) from the activity are positive, $E[\theta] > 0$ and $E[\pi] > 0$. The following proposition gives two necessary conditions for this possibility.

Proposition 3. For ex post review to deter ex ante socially beneficial activities, it is necessary that, under ex post review, the firm benefits from the activity less than the regulator if ex post reversal is costless. That is, it is necessary that

$$\Pi[0, n, g(\theta)] < W[0, n, g(\theta)] \text{ and } E[\pi|\theta \geq -k/n] < E[\theta|\theta \geq -k/n]. \quad (4)$$

Proof. To obtain the necessary condition, note that for a marginal increase in the ex post cost of undoing A, the change in F's ex ante expected profits from A is equal to

$$\frac{\partial \Pi}{\partial k} = \left[\pi\left(-\frac{k}{n}\right) + \frac{k}{n}\right] g\left(-\frac{k}{n}\right) - G\left(-\frac{k}{n}\right), \quad (5)$$

where the first term corresponds to the savings (in terms of profits accrued and break-up costs not spent) associated with the marginal activity that is now not reversed as the reversal cost k increases by an infinitesimal amount, while the second term is equal to the increase in reversal cost for the inframarginal activities (for $\theta < -k/n$) that are still reversed at the higher k. The savings are multiplied by the probability density of the marginal state, while the unit increase in inframarginal reversal costs is multiplied by the probability (equal to the distribution function computed at the marginal state) that the activity is undone in period 1. Comparing equation (5) with (2), we conclude that W decreases faster than Π, $\partial W/\partial k < \partial \Pi/\partial k$, provided that the firm does not want to voluntarily undo the activity, $n\pi(-k/n) > -k$, as guaranteed by Condition 1. Thus, if Π lies above W at $k = 0$, or equivalently, if

$$E[\pi] + n\Pr(\theta \geq 0)E[\pi|\theta \geq 0] \geq E[\theta] + n\Pr(\theta \geq 0)E[\theta|\theta \geq 0],$$

then it is optimal for the firm to undertake the activity ex ante even though the regulator will review it ex post. We conclude that condition (4) is necessary for ex post review to deter an ex ante socially beneficial activity.

To obtain the second necessary condition, note $\Pi[k, n, g(\theta)] < 0, W[k, n, g(\theta)] > 0$, $E[\pi] > 0$, and $E[\theta] > 0$ together imply that $n[1 - G(-k/n)]E[\pi|\theta \geq -k/n] - kG(-k/n) < 0$ and $n[1 - G(-k/n)]E[\theta|\theta \geq -k/n] - kG(-k/n) > 0$. The latter follows from the fact that $W[k, n, g(\theta)] \geq (n+1)E[\theta]$ because W represents welfare from the optimal reversal strategy, while $(n+1)E[\theta]$ is welfare from a strategy of no reversal. This implies that $n[1 - G(-k/n)]E[\theta|\theta \geq -k/n] - kG(-k/n) \geq n[\theta] > 0$. From this, $E[\pi|\theta \geq -k/n] < E[\theta|\theta \geq -k/n]$ follows immediately. *Q.E.D.*

For commitment to forego ex post review to be desirable, there must be positive externalities from the activity under some states. There must be a positive externality in expectation ex ante if reversal costs are zero and there must be an expected positive externality in the event that the activity would not be reversed ex post. This condition is likely to be satisfied when the conflict of interest between R and F is substantial. In this situation, the firm expects to make small profits or even losses when the activity will not be undone by the regulator. Thus, the firm does not want to undertake A because of the cost it expects to bear later to undo the activity when the activity turns out to be ex post socially inefficient. When ex ante the activity is socially beneficial, it is optimal for the regulator to forego the option of ex post review and approve the activity unconditionally. See the Appendix for a simple analytical example of Proposition 3. We also illustrate an example of optimal commitment in our environmental regulation application in Section 6.2.

3.2 Ex ante optimal ex post standard

The first best solution to this conflict between socially optimal control and private incentives is for the regulator to simply compensate the firm for ex post review (of an activity whose ex ante welfare effect is positive) with a subsidy $s \geq kG(-k/n) - n\Pr(\theta \geq -k/n)$ $E[\pi|\theta \geq -k/n] - E[\pi]$. This guarantees that the firm's expected profit from engaging in the activity is non-negative. Alternatively, the regulator could also provide a subsidy of $s/G(-k/n)$ only in the event that it reverses the activity. Either way, such a compensation program would allow the regulator to continue to reverse the activity whenever it is ex post optimal to do so without undermining ex ante incentives to engage in ex ante welfare-improving activities.

If such compensation is infeasible, however, the regulator still may have superior options than foregoing ex post review. Instead, the regulator could commit to a more lenient ex post standard that avoids chilling ex ante socially beneficial activities when $\partial\pi/\partial\theta \leq 0$. If the regulator can commit to not reverse the activity in some welfare-reducing states, it can make the activity ex ante profitable for the firm. If the negative welfare effect of those states is not too large, then the firm now has an incentive to undertake a socially desirable activity while eliminating the reduction in social welfare from the activity in at least some of the states in which the activity is welfare-reducing. One natural approach is to continue to use a single cutoff for reversing the activity but to lower this cutoff so that the activity is allowed in states in which welfare is only slightly negative in order to meet this zero-profit constraint.[4] The following proposition shows that the ex ante optimal ex post standard takes this form.[5]

Proposition 4. Suppose $E[\pi] + n\Pr(\theta \geq -k/n)E[\pi|\theta \geq -k/n] - kG(-k/n) < 0$, $E[\pi] > 0$ and $E[\theta] > 0$. If the regulator can commit to an ex post review policy, then the optimal ex post review policy is to reverse the activity if and only if $\theta \in [\underline{\theta}, \theta^c]$ where $\theta^c < -k/n$ is such that $E[\pi] + n\Pr(\theta \geq \theta^c)E[\pi|\theta \geq \theta^c] - kG(\theta^c) = 0$ under the following conditions: $\pi'(\theta) < 0$, $\pi''(\theta) < 0$ for $\theta < 0$ and $\pi(-k/n) \geq 0$.

Proof. If $\pi(-k/n) < \frac{kG(-k/n)-E[\pi]}{n[1-G(-k/n)]}$, then, because $\pi'(\theta) < 0$, the firm will not undertake the activity unless it is reversed less frequently than is ex post optimal for the regulator. Thus, for the firm to take the action, there must be some interval $[\theta_1, \theta_2]$ with $\theta_2 \leq -k/n$ in which the regulator commits not to reverse it. Let $\pi_{12} = \int_{\theta_1}^{\theta_2} \pi(\theta)g(\theta)$ be the added per period profit the firm earns from approval in this region. Since the firm's profit is concave in the state, for any given π_{12}, we know that the welfare effect of approval in this region, $\int_{\theta_1}^{\theta_2} \theta g(\theta)$, is increasing in θ_2. Thus, the

[4] This is somewhat reminiscent of Mason and Weeds' (2013) result that a more lenient merger policy may stimulate entry.

[5] As the proof shows, the cutoff strategy of reversing the activity if and only if welfare is low enough is optimal because if R ever approves the activity in a region of the state strictly below a region in which it reverses it, if profit is concave in the state, welfare can be increased by shifting that approval region up in a way that keeps the firm's expected profit constant. If profit is sufficiently convex in the state, on the other hand, then it will be optimal to reverse the activity for intermediate θ because we can provide added profit for the firm at the least welfare cost by allowing the activity to continue in a very small region of the worst states for welfare.

regulator can increase the firm's profit by any given amount with the least damage to social welfare by committing not to undo the activity in the region with the highest θ for which it would have otherwise undone the activity. *Q.E.D.*

While both of the last two propositions require some form of commitment ability on the part of the regulator, the commitment required in Proposition 4 is much harder to implement than the one in Proposition 3. It is much easier to verify whether the regulator has reversed a previously approved activity than to determine if the regulator is using a particular ex post standard for reversing the activity. In most cases, the actual standard is not a hard and fast rule but is a reasoned judgment made based on a variety of factors. Thus, even if it is possible to commit not to examine an activity after the regulator has given the activity ex ante approval, it is often practically impossible to commit to a particular standard of ex post review that is not ex post optimal for the regulator.

4. EX POST ASYMMETRIC INFORMATION: DISCIPLINE EFFECT

In this section, we relax the assumption that the state becomes common knowledge in period 1. Instead, we consider what happens if the regulator cannot observe the state, which is instead privately observed by the firm. In addition, the firm can take some action a that might serve to signal the state to the regulator. For example, in the merger application, a would be the post-merger price. A lower price would signal either the merger has greater efficiencies or creates less market power.

Below we derive assumptions on how a affects firm profits that guarantee that this signaling game has a partially-separating equilibrium in which the regulator learns perfectly whether activity A is ex post efficient, that is, whether $\theta \geq -k/n$, in spite of being unable to learn the precise value of the state. In this equilibrium, the regulator's ex post decision is identical to what it would be under complete information.

Formally, we modify the model by assuming that when activity A is undertaken in period 0, F *privately* observes θ in period 1. Furthermore, after learning θ at period 1, F takes an observable action, a, that affects period 1 profit $\pi(a,\theta)$, with $\pi_1(a,\theta) > (<)0$ for $a < (>)\overline{a}(\theta)$. To fix ideas, we also assume that $\pi_{12}(a,\theta) < 0$, so that the action $\overline{a}(\theta)$ that maximizes period 1 profit is decreasing in the state, θ.[6] (The analysis would be similar for $\pi_{12}(a,\theta) > 0$, so that the action $\overline{a}(\theta)$ would be increasing in the state.)[7] As explained above, we assume that R observes this action, a, but does not observe the state, θ.

We also allow for action a to affect social welfare. To account for this in our model, we slightly refine our prior definition of social welfare. Social welfare is given by $\omega(a,\theta)$ with $\omega_2 > 0$ and $\omega(\overline{a}(\theta),\theta) = \theta$. This means that, as we assumed above, social welfare is given by θ if F chooses its profit-maximizing action in period 1. Since θ was commonly observable in period 1 in our baseline model, there was no reason for F to choose any different action. Thus, this extension is consistent with our definition of θ as social welfare in the previous sections.

It is convenient to break up the analysis into three steps. First, Section 4.1 analyzes the firm's best response when the regulator commits to an ex post regulatory policy. Second,

[6] In period 2, F will have no reason to deviate from its unconstrained optimal action, $\overline{a}(\theta)$, resulting in payoff $n\pi(\overline{a}(\theta),\theta)$, which is what we denote by $n\pi(\theta)$ as in the baseline model.

[7] All these conditions are satisfied in the merger application in Section 6.1 when $\pi(\theta)$ is increasing. a is the period 1 price, profit initially increases (and beyond some level decreases) in price, and increasing the price raises profits less (or reduces profits more) the larger is θ, that is, the more competitive the market is or the greater the efficiencies from the merger. We explicitly illustrate the discipline effect in our discussion of the environmental permit application with Cournot competition.

Section 4.2 establishes conditions for a partially-separating equilibrium to exist. Third, Section 4.3 establishes conditions under which the regulator would like to commit to a stricter ex post regulatory policy.

4.1 Firm behavior

Let $\tilde{\pi}(\theta) = \pi(\overline{a}(\theta), \theta)$ be F's one-period profit in state θ when the one-period profit-maximizing action a is chosen (this corresponds to $\pi(\theta)$ in our baseline model). Assuming the regulator adopts a cutoff policy in terms of the action taken by the firm in period 1, we obtain the following result.

Lemma 1. Suppose that the regulator undoes activity A in period 1 if and only if $a > \hat{a}$. If $\pi_2(a, \theta) > 0$ and $n > 1$, then the firm chooses $\overline{a}(\theta)$ for values of θ such that $\tilde{\pi}(\theta) - k > \pi(\hat{a}, \theta) + n\tilde{\pi}(\theta)$ and for values of θ such that $\overline{a}(\theta) < \hat{a}$. The firm chooses $\hat{a}$ otherwise. That is, if θ_1 and θ_2 are defined by $\tilde{\pi}(\theta_1) - k = \pi(\hat{a}, \theta_1) + n\tilde{\pi}(\theta_1)$ and $\overline{a}(\theta_2) = \hat{a}$, the firm chooses $\overline{a}(\theta)$ for $\theta \subset [\underline{\theta}, \theta_1] \cup (\theta_2, \overline{\theta})$ and chooses $\hat{a}$ for $\theta \subset (\theta_1, \theta_2]$.

The lemma says that if social welfare and firm profit are both increasing in the state and the regulator adopts a policy of undoing A if and only if a is sufficiently large, then if the state is sufficiently low, the firm gives up on convincing the regulator otherwise and simply chooses its one-period profit-maximizing outcome and the regulator undoes A. In this case (as illustrated in Figure 3), the profitability of A is low enough that it is not worth sacrificing current profits to induce the regulator not to undo A. For a middle range of states, however, the firm chooses the maximum action for which the regulator will not undo A. Even though this action is smaller than the one-period profit-maximizing option, the firm in this range of states finds it profitable because the period 1 loss in profit is small enough relative to the future gain of continuing the activity. Then, for very high states, the firm can choose its one-period profit-maximizing action and still have the regulator not undo A.

It is worth noting that the condition that the firm's profit is increasing in θ is not innocuous.[8] If the firm values the activity more when it is also more socially beneficial, it has the incentive to distort its behavior to signal that the state is larger. If firm profits are greater when the state is smaller, however, then such signaling is not possible because low θ types have a larger benefit from ex post approval and, therefore, more reason to distort behavior to be approved ex post.

4.2 Equilibrium analysis

Given the firm's best response, we now examine the regulator's ex post optimal strategy. In the best equilibrium for R, it would set $\hat{a}$ so that $\theta_1 = -k/n$, the state for which the social loss from allowing A to continue exactly equals the social cost of undoing the action.[9] That is, R sets $\hat{a} = \hat{a}^*$, where $\hat{a}^*$ satisfies $\pi(\hat{a}^*, -k/n) + (n-1)\tilde{\pi}(-k/n) = -k$. This will generate the semi-separating equilibrium in which the regulator reverses the activity if and only if

[8] For merger review, this condition will hold if θ measures efficiencies from a merger, but does not hold if θ represents the amount of competition after a merger. For applications to environmental economics, the condition is more easily satisfied since the conflict of interest between the firm and the regulator is generally one of different weighting of interests rather than having directly opposing interests as is the case in antitrust regulation with respect to market power. For instance, if θ represents the ability of a new plant to minimize environmental damage or the ease of extracting oil in an environmentally sensitive area, both profit and social welfare should move in the same direction.

[9] This is the best equilibrium for R ex post. As we discuss below, it is possible that the regulator would want to commit to an even lower level of $\hat{a}$.

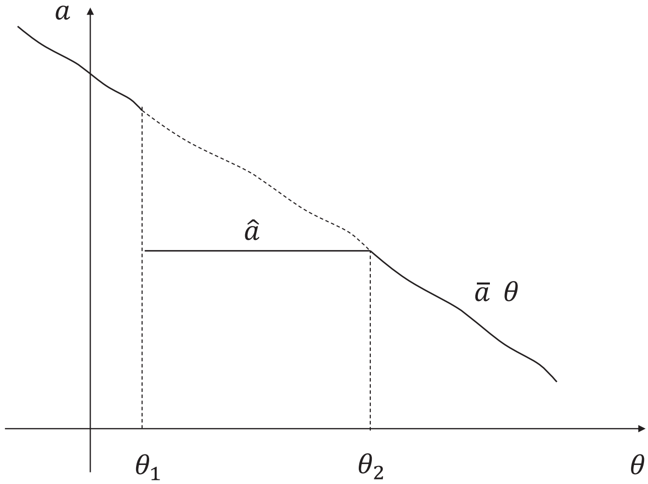

Figure 3. Signaling equilibrium.

$\theta < -k/n$ because those are the only states in which the firm chooses $a > \hat{a}^*$. This replicates the full information outcome.

This equilibrium, however, is not unique. In all other equilibria, activity A is allowed to continue in some states in which undoing it would increase social welfare. This occurs because there is partial pooling across states for which allowing A to continue increases social welfare with states for which doing so decreases social welfare. In the rest of this section, however, we focus on the welfare-maximizing equilibrium. Appendix A.2 analyzes the other possible equilibria in more detail.

In a signaling equilibrium, the firm's action in period 1 is the same as in the full-information case except when $\theta \in (\theta_1, \theta_2]$. In the welfare-maximizing equilibrium, the regulator's ex post decision is also the same as in the full information case. Thus, the effect of ex post asymmetric information (at least in the welfare-maximizing equilibrium) relative to full-information depends on the welfare effect of the firm's choosing $\hat{a}^*$ rather than $\overline{a}(\theta)$. If $\omega_1(a, \theta) < 0$ for $a \in [\hat{a}, \overline{a}(\theta)]$, then ex post asymmetric information in the welfare-maximizing equilibrium will lead to greater social welfare than in the full-information case.[10] Moreover, ex post review can have the additional advantage of disciplining the firm to choose actions closer to the social optimum in period 1, thus increasing the benefit of ex post review relative to ex ante review.

4.3 Commitment

Lastly, note that even relative to the ex post welfare-maximizing equilibrium, there are circumstances in which the regulator could benefit from the ability to commit to a standard different from $\hat{a}^*$. If social welfare is greater in the ex post welfare-maximizing equilibrium reported in the previous section than in the case with symmetric information, it is because having the firm choose $\hat{a}^*$ rather than $\overline{a}(\theta)$ for $\theta \in (\theta_1, \theta_2]$ increases social welfare. Then, a reduction in the action a in this region might further increase social welfare, which will be the case provided that $\omega_1(a, \theta) < 0$ at $a = \hat{a}^*$. If R can commit to choosing $\hat{a} < \hat{a}^*$, the

[10] Similarly, if $\pi_{12}(a, \theta) > 0$, the social welfare would be greater than in the full information equilibrium if $\omega_1(a, \theta) > 0$ for $a \in [\overline{a}(\theta), \hat{a}]$ because the firm would be distorting a upward to $\hat{a}$ to signal the state.

boundaries θ_1 and θ_2 will be affected along with the action a that F will choose in this region. Explicitly writing θ_1 and θ_2 as a function of $\hat{a}$, we can analyze the effect of a small reduction in $\hat{a}$ from $\hat{a}^*$ to $\hat{a}^{**}$.

Proposition 5. Suppose that $\omega_1(a, \theta) < 0$ for $a \in [\hat{a}^*, \overline{a}(\theta)]$ and $\pi_2(a, \theta) > 0$. If the regulator can commit to the set of actions for which it will reverse activity A, then it sets this cutoff at $\hat{a}^{**} < \hat{a}^*$ if $g(\theta_1(\hat{a}^*))$ is sufficiently small.

A small reduction in $\hat{a}$ shifts up the region of states (i.e., shifts up θ_1 and θ_2) for which the firm lowers its action in order to signal a higher state.[11] Hence, the change in $\hat{a}$ from $\hat{a}^*$ to $\hat{a}^{**}$ brings about three effects on social welfare. First, for states (between θ_1 and θ_2) in which the firm was going to choose its action to just meet the standard, $\hat{a}$, this action is now reduced from $\hat{a}^*$ to $\hat{a}^{**}$. This reduction in the action increases social welfare under our working assumption that lower actions increase welfare in this range. Second, because the region of states in which the firm lowers its action to signal a higher state shifts up, there is a region (between the old θ_2 and the new, higher, θ_2) in which the firm is now lowering its action to signal a higher state when it otherwise would not need to do so. This reduction in the action also increases social welfare. Third, there is now a region of states (between the old θ_1 and the new, higher, θ_1) in which the firm increases its action because it no longer tries to signal that the state is sufficiently high. This increase in the action reduces social welfare. According to Proposition 5, if the probability of being in this third region is sufficiently small, then the regulator would benefit from being able to commit to lowering the cutoff action $\hat{a}$ below the ex post optimal level because doing so would induce a socially desirable change in the firm's period 1 action.

5. ENDOGENOUS REVERSIBILITY COST: MANIPULATION EFFECT

In this section, we relax the assumption that k, the ex post cost of undoing A, is exogenous. Instead, we allow the firm to choose this cost in period 0 after the regulator has decided not to prohibit A ex ante. As we will see, ex ante review and possibly commitment to ex ante review become more desirable when the firm can manipulate this cost. To endogenize k, suppose that first-period profits are given by $\pi(k, \theta)$, where $\pi_1(k, \theta) > (<) 0$ for $k < (>) \overline{k}$, $\pi_{11} < 0$, and $\pi_{12}(k, \theta) = 0$. That is, there is a level of k, $\overline{k}$, that maximizes period 1 profit. For simplicity, we assume this level of reversibility is independent of the state θ. To isolate the manipulation problem from the regulator's concern for ex post social welfare, we further assume there are no direct externalities associated with k, in the sense that changes in k only affect period 1 social welfare through their effect on profits. However, changes in k may also impact social welfare either through their effect on R's reversal decision or through the cost of undoing A should R decide to do so. Thus, to be consistent with our interpretation of θ as a measure of social welfare, we define social welfare in period 1 by $\tilde{\omega}(k, \theta)$ with $\tilde{\omega}_2 = 1$, $\tilde{\omega}_1 = \pi_1(k, \theta)$, and $\tilde{\omega}(\overline{k}, \theta) = \theta$. This means that social welfare is given by θ if F chooses its profit-maximizing k in period 1, as we assumed in the baseline model. Thus, our description of θ measuring social welfare in the previous sections is consistent with our definition here. Period 2 payoffs are as in the baseline model.[12]

[11] In the merger context, the efficiencies have to be slightly larger for the firm to charge the newly lowered price necessary to signal large enough efficiencies so as to avoid reversal of the deal.

[12] When we write π as a function of just one argument, we are referring to profits accrued at period 2. Whereas, if π is a function of two arguments, we are referring to profits at period 1.

Assuming that R cannot commit ex ante, so that its reversal standard remains at $-k/n$, we obtain the following result.

Proposition 6. Assume that π_{11} is sufficiently negative to guarantee that $\partial^2\Pi/\partial k^2 < 0$.[13] If the firm can choose k in period 0 and R chooses whether or not to undo activity A to maximize ex post social welfare, then the firm's choice of k exceeds the socially optimal level.

Proposition 6 potentially gives another reason for ex ante control. Without the ability to commit to ex ante control, the firm will choose to make it more costly than socially optimal to undo the activity ex post. That said, note that the proposition does not establish that the privately optimal k necessarily exceeds $\bar{k}$. That is, prior to this section, we implicitly assumed a fixed k. One could think of this as assuming that k is fixed at $\bar{k}$, that is, the firm has to choose the level of reversibility costs that maximize current profits. If so, relaxing this assumption does not necessarily suggest that ex ante control is more likely to be superior relative to the situation where F has no flexibility to adjust k.

For illustration, note that F's distortion of k away from $\bar{k}$, the period 1 profit-maximizing level, is due to both a desire to reduce the probability of reversal (which pushes towards a larger k) and a desire to minimize the losses when reversal occurs (which pushes towards a smaller k). F will choose a k that exceeds the period 1 profit-maximizing level if and only if its expected future gains from a reduced probability of R reversing A exceed the increased loss when A is still reversed. Interestingly, F may be more likely to choose $k > \bar{k}$ if $\bar{k}$ is already quite large so that the probability that the state is low enough to warrant reversal of A is small (because reversals are so socially costly). That said, if very low states are quite unlikely (if, e.g. g is single humped and close to zero at $\underline{\theta}$), then this may not be the case since increasing k from this already high level may not reduce the probability that R reverses A very much.

Note, the manipulation effect assumes that R does not have unlimited commitment power. If R could commit to reverse A whenever k differed from the socially optimal level, then F would not deviate from this level. However, R might not have this commitment power, even if it could commit to only ex ante review, because it might not be able to describe the socially optimal k in advance. Alternatively, if commitment is possible due to reputational effect, it might be too difficult for outsiders to determine if k differed from the socially optimal level, so that the reputation loss from not reversing would be minimal.

6. APPROVAL WITH CONDITIONS

In our baseline analysis, we assume the regulator can either approve or prohibit the activity, in either period 0 or period 1. In many contexts, however, regulators choose to allow the activity subject to some conditions. For example, mergers can be allowed if the merging parties divest some assets. Energy exploration can be allowed if certain environmental precautions,

[13] Note that

$$\frac{\partial^2\Pi}{\partial k^2} = \pi_{11}(k,\theta) + \frac{1}{n}\left\{\left[2 - \pi'\left(-\frac{k}{n}\right)\right]g\left(-\frac{k}{n}\right) - \left[\pi\left(-\frac{k}{n}\right) + \frac{k}{n}\right]g'\left(-\frac{k}{n}\right)\right\}.$$

If g is single humped, then ensuring $\partial^2\Pi/\partial k^2 < 0$ is most difficult at $g'(-k/n) = 0$. Thus, the density of the distribution cannot be too large at its peak relative to the concavity of profit in k. Of course, for a sufficiently large n, this condition should be relatively easy to satisfy.

limitations, or remedial measures are employed. Loans can be offered if certain disclosures are made and under certain limits on the repayment terms.

We can incorporate conditions into our model in a variety of ways. The simplest way is to assume that these conditions affect welfare and profit in a fixed and deterministic way. In that case, we can simply think of the activity as the activity with the socially optimal conditions imposed by the regulator. With this interpretation, our analysis is completely unaffected.

Alternatively, we could stipulate that the optimal conditions vary with the state, while also assuming that allowing the activity with conditions does not affect the ability to learn the state. In that case, our analysis also applies with little change. The only difference is that one might allow the activity initially subject to some conditions and then change those conditions in period 1 based on the realization of the state θ.

This would potentially change the need for ex ante commitment. Commitment would become less valuable if ex post blocking of the activity were replaced with conditions under which the activity was still profitable. If instead the regulator were more likely to impose welfare-increasing but profit-reducing conditions ex post, the value of ex ante commitment would be enhanced.

Lastly, one could imagine that introducing conditions could affect the ability to learn the state. The analysis would depend on how the conditions impacted the learning process. For example, if imposing conditions significantly reduced the ability to learn the state, approving with conditions ex ante would be similar to blocking the activity ex ante. In this case, our analysis remains valid once the baseline level of welfare is modified.

7. APPLICATIONS

In this section, we will illustrate some of the key elements of our results in two applications of the model, merger review and environmental permit regulation. Of course, the model's insight could apply to other types of approval regulation when there is uncertainty about the benefits of the activity. Additional applications include vertical merger review where there is not only a foreclosure concern but also uncertainty about the degree to which the merger would facilitate welfare-improving investments, and new drug approvals in which the uncertainty is about either the therapeutic benefits or safety risks of the drug. Perhaps the most straightforward application is to consumer product regulation in which the product is genuinely beneficial to sophisticated consumers but harmful to naive consumers, who wrongly believe it to be beneficial. If the size of one of these two groups is uncertain, our results apply directly. If firms can lure naive consumers into purchases through marketing, a discipline effect could operate in which ex post review induces less of this welfare-reducing marketing.

7.1 Merger review

In the merger context, commitment may be desirable if the merger is likely to generate substantial efficiencies, but those efficiencies will mostly be passed on to consumers when the merger results in a small increase in market power. If the merger substantially increases market power, the regulator will want to undo the merger. Then, the merger may not generate enough profits (when the regulator lets it stand) to compensate the firm for the costs of having to undo the merger. Ex ante, however, the merger may be socially beneficial if the efficiency benefits to consumers in the low market power states are quite large and these states are quite likely.

Focusing on our leading application to merger control, this section illustrates how the payoffs functions for F and R can be derived from a fully-fledged industrial organization model.

We consider a merger to monopoly between two symmetric firms $n = 1, 2$ offering differentiated products and competing simultaneously in prices. Denoting the price and the quantity of product i by p_i and q_i, assume that demand is given by Shubik and Levitan's (1980) linear demand system, with $q_i =: (2v - 2p_i + \lambda(p_j - p_i))/4$ for $i, j \in \{1, 2\}$ with $j \neq i$, where $\lambda \in [0, \infty)$ measures the degree of substitutability between the two products. Before the merger the marginal cost of production for each firm is c. The merger reduces the marginal cost to $c - \gamma$, where $\gamma \in [0, c]$ measures the extent of the synergy created by the merger.

In the Bertrand–Nash equilibrium resulting in the baseline scenario in which the firms do not merge, equilibrium profits for each firm and total surplus in equilibrium are equal to

$$\pi_1 = \pi_2 = \frac{\lambda + 2}{(\lambda + 4)^2}(v - c)^2 \text{ and } TS = \frac{(\lambda + 2)(\lambda + 6)}{2(\lambda + 4)^2}(v - c)^2,$$

where $\pi_i = (p_i - c)q_i$, $CS = (v - p_1)q_1 + (v - p_2)q_2 - [1/(1 + \lambda)][q_1^2 + q_2^2 + (\lambda/2)q_1 q_2]$, and $TS = \pi_1 + \pi_2 + CS$. Once the firms merge, profits and total surplus are equal to

$$\pi_{1\&2}^A = \frac{1}{4}(v - c)^2 + \frac{1}{4}\gamma[2(v - c) + \gamma] \text{ and } TS^A = \frac{2\lambda + 3}{8(\lambda + 1)}(v - c)^2 + \frac{2\lambda + 3}{8(\lambda + 1)}\gamma[2(v - c) + \gamma].$$

We consider two alternative cases, depending on whether the uncertainty is about cost synergy, γ, or product substitutability, λ. First, we show that when the synergy is uncertain, the change in equilibrium profits induced by the merger increases in the change in total surplus achieved in equilibrium. Second, when instead product substitutability is uncertain, the change in equilibrium profits induced by the merger decreases in the change in total surplus.

7.1.1 Uncertain synergy: increasing $\pi(\theta)$

For any given level of substitutability, $\lambda \in [0, \infty)$, we now characterize the impact of the synergy, γ, on the change in profits and total surplus generated by the merger. Note that the change in profits for the merging firms, $\pi_{1\&2}^A - \pi_1 - \pi_2$, and the change in total surplus from the merger, $TS^A - TS$, are both increasing functions of γ. This property follows immediately from the observation that γ affects positively post-merger profits as well as social surplus, but does not affect pre-merger profits and social surplus. Thus, once we define the state as $\theta := TS^A(\gamma) - TS(\gamma)$, the differential profits from the merger, $\pi(\theta) := \pi_{1\&2}^A - \pi_1 - \pi_2$, are increasing in θ.

The discipline effect can operate in this application if a represents the post-merger price (where we can constrain the two prices to be equal because the products are symmetric), since $\hat{a}^* \leq \overline{a}(\theta)$ for $\theta \in (\theta_1, \theta_2]$ (with equality only at $\theta = \theta_2$). Notice that $\omega_1(a, \theta) < 0$ for $a \in [\hat{a}, \overline{a}(\theta)]$ whenever $\hat{a}^*$ exceeds marginal cost at θ_1, so that ex post asymmetric information simply brings price closer to marginal cost. This is possible since the firm has market power after the merger, so that its profit-maximizing price exceeds marginal cost. In this case, the signaling distortion that reduces profit for the firm is a social benefit, making the case for ex post review stronger when there is ex post asymmetric information. Of course, if the signaling equilibrium is something other than the welfare-maximizing one, then not only will the pooling action be larger, but the regulator's decision will be imperfect, making the case for ex post review weaker.

7.1.2 Uncertain substitutability: decreasing $\pi(\theta)$

Next, holding constant the synergy, $\gamma \in [0, c]$, we characterize the impact of the substitutability, λ, on the change in profits and total surplus generated by the merger. The change in total surplus from the merger, $TS^A - TS$, is a decreasing function of λ. To see this, note that the post-merger level of social surplus TS^A decreases in λ in the absence of the synergy, further decreases in λ when the synergy is strictly positive (the more so, the stronger the synergy), and the pre-merger level of social surplus TS increases in λ. Intuitively, the more closely substitutable the products, the bigger the deadweight loss resulting from market power when the firms merge and act as a joint monopolist that controls both prices. On the other hand, the more closely substitutable the products, the stronger the competition between the firms when they are separate, and thus the smaller the resulting deadweight loss. By the same logic, the change in profits for the merging firms, $\pi^A_{1\&2} - \pi_1 - \pi_2$, is an increasing function of λ, because $\pi^A_{1\&2}$ is constant in λ while π_1 and $\pi_2 = \pi_1$ decrease in λ. Once we define the state as $\theta := TS^A(\lambda) - TS(\lambda)$, we conclude that $\pi(\theta) := \pi^A_{1\&2} - \pi_1 - \pi_2$ is decreasing in θ. In this case, the preferences of the firm and the regulator go in opposite directions: an increase in the product substitutability between the products enhances the merged firm's ability to leverage its market power, and thus makes the merger more profitable but less socially desirable.

It is worth noting that this application satisfies the assumptions that $\pi'(\theta) < 0$ and $\pi''(\theta) < 0$ for $\theta < 0$ in Proposition 4 if and only if $\lambda > 2$. This follows because $\pi^A_{1\&2} - \pi_1 - \pi_2$ is concave in λ if and only if $\lambda > 2$ and $TS^A(\lambda) - TS(\lambda)$ is always convex in λ. In other words, as long as the products are close enough substitutes, so that the merger substantially increases market power, further increases in market power from the merger have a diminishing benefit for the firm. On the other hand, increases in market power when the market was close to being perfectly competitive have increasing returns. So, if the merger is between close enough substitutes, then Proposition 4 describes how to determine the ex ante optimal ex post standard.

7.2 Environmental permit regulation

Turning to environmental regulation, suppose the government has to approve the construction of a new plant in an environmentally sensitive area or the use of a particular type of production technology or the extraction of a natural resource where such extraction might cause environmental damage. To fix ideas, say a firm wants to undertake an activity in an environmentally sensitive location (e.g., drilling oil in the Arctic or building a dam). The project generates profit, consumer surplus, and environmental harm.

One possibility is that the environmental harm is known, $L > 0$, but there is uncertainty about the availability of alternative energy sources, z. Let $\pi(z)$ and $\theta(z) = \pi(z) + \pi^A(z) + CS(z) - L$ be the firm's profit from the project and the net welfare effect of the project respectively, where $\pi^A(z)$ is the profit of existing energy producers and $CS(z)$ is the increase in consumer surplus from the activity. To be concrete, say the project generates a quantity x and price of energy is given by $1 - (x + z)$ and there is Cournot competition. If we assume zero marginal cost of energy production, then

$$\pi(z) = x(1 - x - z).$$

Thus, profit is decreasing in z. Total welfare is given by

$$(x + z)^2/2 - z^2/2 + (x + z)(1 - x - z) - z(1 - z) - L. \tag{6}$$

The first two terms are the change in consumer surplus from the added supply, the second two terms are the change in industry profits and the last term is the environmental loss from the activity. It is easy to see that $\pi(z)$ and $\theta(z)$ are decreasing functions of z, thus $\pi(\theta)$ is increasing.

If the relevant uncertainty is about alternative energy supplies, the discipline effect can make ex post review more desirable. To see this, imagine that, instead of the energy generated by the project being fixed, it is a choice variable for the firm, so that the action, a, is just the choice of x. The profit-maximizing choice is $x = 1/2 - z$, generating a maximum profit of $(1 - z^2)/4$. Simplifying equation (6), we can write per period social welfare from the project as

$$x(2 - x - 2z)/2 - L. \tag{7}$$

Thus, if F chooses the profit-maximizing x, the social welfare gain from the activity is given by $3(1 - z)^2/8 - L$, which means the activity is welfare-increasing for $z < 1 - 2\sqrt{2L/3}$.

Using the condition in Lemma 1, we can derive $\hat{x}^*$, the value of x for which the firm is indifferent between choosing this x and having the regulator not reverse the activity versus choosing its profit-maximizing level given that the regulator reverses the activity when $z = 1 - 2\sqrt{2L/3}$. This is

$$\hat{x}^* = \sqrt{2L} + \sqrt{3k - 2L(n-2) + \sqrt{6L}(n-1)}.$$

Notice that $\hat{x}^*$ is increasing in k and increasing in n for $L < 3/8$ (if $L > 3/8$ then the activity is never welfare-increasing). At $k = 0$ and $n = 1$, $\hat{x}^* = 2\sqrt{2L/3}$, the socially optimal x for $z = 1 - 2\sqrt{2L/3}$. Thus, if $k = 0$ and $n = 1$, if the regulator reverses the activity ex post if and only if $x < \hat{x}^*$, then the activity continues after period one if and only if it is welfare-increasing while also generating weakly greater first-period welfare. For larger k or n, the $\hat{x}^*$ which generates the complete information regulatory outcome will produce socially excessive first-period output for at least some z for which the activity is welfare-increasing at the profit-maximizing x.

Alternatively, it could be that the degree of environmental harm, L, is uncertain. In this case, π would likely be independent of L, but $\theta(L)$ would be decreasing. In that case, π would also be independent of θ. We can illustrate the advantages of committing not to review ex post if we imagine that $L = \{0, L_h\}$, letting q be the probability that $L = L_h$. With per period social welfare from the activity given by (7), it is ex post optimal to reverse the activity when $L = L_h$ if and only if $k < n[L_h - x(2 - x - 2z)/2]$. If this holds, the firm's expected profit from ex post review is $n(1 - q)x(1 - x - z) - qk$. So if $k > n(1 - q)x(1 - x - z)/q$, then the firm will not undertake the activity with ex post review. The activity is ex ante welfare increasing, however, if and only if $x(2 - x - 2z)/2 - qL_h > 0$ or $L_h < x(2 - x - 2z)/(2q)$. In other words, commitment to not review ex post is optimal if and only if

$$x(2 - x - 2z)/2 + k/n < L_h < x(2 - x - 2z)/(2q) \text{ and } n(1 - q)x(1 - x - z)/q < k.$$

Because $x(2 - x - 2z)/(2q) > x(2 - x - 2z)/2 + k/n$ for $k = (1 - q)x(1 - x - z)/q$, there exist a range of possible L_h and k for which commitment to ex ante review is optimal.

8. LITERATURE

At a broad level, our modeling approach follows the literature on optimal delegation without transfers, initiated by Holmström (1984) and advanced by Amador et al. (2006), Alonso and Matouschek (2008), and Armstrong and Vickers (2010) among others. Our contribution is to embed a special version of the delegation problem (with complete information and a binary action set) into a simple dynamic collective experimentation problem (with two periods). While in the optimal delegation literature the regulator is free to constrain the action set to any subset of the set of feasible actions, in our setting the regulator can prohibit the activity but cannot force it.

The value for a principal from committing to an ex post sub-optimal standard to obtain a beneficial change in the agent's behavior is familiar from the contracting literature—see, for example, Taylor and Yildirim's (2011) analysis of the commitment value of blind reviews. To the best of our knowledge, however, the tradeoff we identify here between ex ante and ex post approval has not been generally analyzed before. In their informal discussion of ex ante and ex post modes of government regulation, Viscusi et al. (1995: 785–6) notice how liability may chill research incentives. In his broad discussion of the timing of government oversight, Rey (2003: Section 4.2) identifies informally, among other factors, the value of flexibility to adapt to ex post circumstances and the value of commitment not to exploit the firm ex post. These verbal discussions anticipate the chilling effect characterized in our formal model. In the context of a binary version of the baseline model we present here, Ottaviani and Wickelgren (2011) apply to merger review some (but not all) of the effects we identify in this more general setting. This article provides the precise conditions for analyzing the tradeoffs.

Besanko and Spulber (1993) analyze the strategic effects of the merger approval process when the merging firms have private information about the efficiency consequences of the merger. They model how commitment to a consumer welfare standard influences the merger filing decision in a way that increases social welfare.[14] Given our different focus on learning about the effect of the merger, we instead analyze a setting in which the firm has no preexisting private information. Ex post information arrives only if the merger takes place—in our dynamic setting we then compare ex ante with ex post review.[15]

Shavell (1984) compares ex ante and ex post regulation but focuses on a different tradeoff from the one we identify here. He obtains an informational advantage for ex post tort liability over ex ante safety regulation because the ex post regime induces injurers to take optimal ex ante precautions on the basis of their private information.[16] In his model, the only reason the ex post regime does not achieve the first best is that ex post liability is insufficient due to judgment-proof injurers and plaintiffs who do not always sue.[17]

Schwartzstein and Shleifer (2013) model the tradeoff between ex ante regulation and ex post liability in a setting with positive externalities from the activity. In their model, there is

[14] We also refer to Farrell and Katz (2006) for a broad discussion of the incentive effects of different welfare standards in merger policy.

[15] Barros (2003) analyzes the European Union's shift from a system requiring prior notification for horizontal agreements to a system of ex post control. He focuses on the effect of this change on the restrictiveness of the agreements firms propose. Katz and Shelanski (2007) critique the US antitrust agencies handling of uncertainty in merger review. Salop (2016) suggests reviewing past consent decrees to determine if they are sufficient to eliminate anti-competitive effects, while Patel (2020) argues for more ex post merger review.

[16] Daripa and Varotto (2010) contrast ex ante with ex post regulation of bank capital. As in Shavell's model, the advantage of ex post regulation is that it allows the bank manager to use her private information about risk. On the other hand, ex post regulation is more vulnerable to unknown managerial risk aversion.

[17] Kolstad et al. (1990) also analyze ex ante safety regulation and ex post tort liability. In their model, uncertainty about how a court will determine negligence can lead to inefficiencies in ex post tort liability, which can be corrected with appropriate ex ante regulation.

no learning (regulation is only efficient when regulators are more accurate than courts), but the regulated firms have private information as to their safety. Because there is only one period of activity in their model, learning is not relevant. But, they do find an effect reminiscent of our commitment effect in that ex ante regulation can reduce the inefficient deterrence of socially desirable activity. Bose et al. (2016) also examine the tradeoff between ex ante price regulation and ex post antitrust liability in a more detailed industrial organization model without learning. Kotowski et al. (2014) model ex ante rules versus ex post standards with private information on compliance costs. In their model, ex post regulation allows inferences from firm actions but there is no direct information from seeing the result of the action. They do find an effect related to our discipline effect except that it goes in the opposite direction—firms with low compliance costs don't comply as much under ex post standards to influence the regulator to believe that compliance costs are high.

Our model contributes to the understanding of the option value of delay in strategic environments when uncertainty is resolved over time. We depart from the classic literature on real options literature (see Arrow and Fisher 1974; Henry 1974) in two important ways. First, in our model information arrives only if the activity is undertaken, while in the real options literature information arrives exogenously. Thus, there is active experimentation rather than passive learning. Second, our two decision-makers (firm and regulator) interact strategically, while the real options literature mostly focuses on nonstrategic environments.[18]

More closely, our model contributes to the literature on collective experimentation, spearheaded by Strulovici's (2010) analysis in a voting context. Our setting, instead, features a firm biased toward the activity and a regulator who can block the activity but cannot force the firm to undertake it; in the baseline specification, the two players learn the state of the world perfectly but only if the activity is undertaken, which requires unanimity—undoing the activity at the point is costly.[19] In a complementary contribution, Henry et al. (2022) analyze a two-phase continuous-time model building on Henry and Ottaviani's (2019) collective experimentation setting à la Wald (1945). In a first phase, noisy information is collected through costly experimentation, resulting in product adoption or abandoning experimentation altogether. In a second phase, following adoption, monitoring performance provides additional noisy information, possibly leading to product withdrawal if sufficiently bad news is revealed. In the two-period model proposed here, instead, information revelation is perfect once the activity is undertaken, but unscrambling the eggs is costly.

Kwoka and Valletti (2021) overview a number of historical cases of successful breakups that eventually led to lower prices and more innovation, thus challenging the view that breakups of previously approved mergers are excessively costly and impractical. As predicted by our model, breakups can be an effective policy tool in sectors like tech in which the impact of mergers is difficult to predict ex ante.

9. CONCLUSION

Our analysis of the tradeoffs inherent in ex ante and ex post regulation provides some insight into how regulators should use the option of ex post enforcement along with ex ante enforcement. We show that the option to prohibit some activity after the regulator becomes better

[18] The strategic interaction that has been considered in the real options literature is very different from ours. For example, in Weeds' (2002) analysis of strategic preemption, the players are firms that compete to invest after learning has taken place, whereas in our model the regulator and the firm must agree ex ante for the activity to be undertaken.

[19] Note that collective experimentation is fundamentally different from strategic experimentations (as analyzed, e.g. by Bolton and Harris 1999) in which agents experiment independently from each other, rather than collectively deciding whether to experiment.

informed about its consequences should often induce the regulator to become more lenient in its ex ante decision. The magnitude of this leniency is greater the smaller is the cost of undoing an activity after it has been undertaken (the cost of “unscrambling the eggs” in the merger context), the more uncertain are the effects of the merger, and the longer the activity will continue to have effects in the market after the ex post decision.

Our analysis also uncovers the benefits of allowing the regulator to commit not to undertake ex post review when there is sufficient conflict of interest between the regulator and the firm proposing the activity. This commitment is valuable in situations in which the activity is socially beneficial in expectation, but the states in which it increases social welfare are also (largely) the states in which it reduces firm profit. The firm may then be discouraged to take on such a welfare-increasing activity when the regulator is unable to commit not to review the activity ex post. In these circumstances, the regulator also benefits from committing to a more lenient ex post standard.

We also demonstrate that the presence of ex post asymmetric information (according to which the firm learns the true effect of the activity but the regulator does not) often does not reduce the benefit of ex post review and may actually increase it. In the merger context, if the nature of the uncertainty concerns the efficiencies the merger will generate, then if the firm has private information regarding these effects after merging, it will want to signal that the efficiencies are large by charging low prices. If prices exceed marginal cost because of market power, this signaling distortion may actually increase social welfare.

Ex post review, however, might create a perverse incentive for the firm to increase strategically the cost of reversing the activity so as to essentially force the regulator not to undo the activity. Note that all three effects identified here (chilling, discipline, and manipulation) result in a reduction of the number of ex post reversals, consistent with the rare occurrence of reversals in the application to merger control.

Our model assumes the regulator puts equal weight on the welfare of the firm and the rest of society (consumers in the competition applications, the broader population in the environmental or health applications). While in many regulatory contexts, this is a reasonable assumption, in most countries competition regulators follow a consumer welfare standard, putting zero weight on the welfare of the firm(s).

Most of our results are robust to this alternative welfare function. Our model assumed the full reversal cost, k, is borne by the firm. While this assumption does not matter much under our total welfare function, it makes a more significant difference under a consumer welfare function. If we retain this assumption, the cutoff for ex post reversal under consumer welfare standard changes from $\theta = -k/n$ to $\theta = 0$. In addition, in the absence of commitment, ex ante regulation is almost never optimal. On the other hand, commitment will be more important, especially if profit is decreasing in the state, because there will be more ex post reversals that result in losses for the firm. The discipline effect will continue to operate similarly, but there will be no possibility of manipulation if the regulator puts no weight on firm profits.

If, instead, we assume that part of the reversal cost is external to the firm, and thus counts as part of the regulator’s welfare function, our results are all qualitatively similar. Commitment will continue to be somewhat more important. The manipulation effect will continue to exist but the firm will focus on increasing the external cost of reversal.

Outside of merger review, our analysis applies to other areas of competition policy, such as in the regulation of cooperation agreements among competitors as well as vertical or exclusive dealing contracts where there are potentially both pro- and anti-competitive effects.[20]

[20] Consider, for example, the framework for regulating agreements among competitors. Through Council Regulation (EC) 1/2003 (so-called Modernization Regulation), the European Union replaced an old ex ante control system (originally

More broadly, the tradeoffs uncovered in our simple model arise in other areas in which a regulator has to decide whether to allow or prohibit some activity that may or may not be harmful to society and for which uncertainty is reduced if the activity is allowed for some time and its effects are observed. While we applied our model mostly to competition policy, the analysis is relevant also for environmental and health regulation as well as for professional licensing and urban planning. A multi-period extension of the model could be developed to analyze the optimal multi-phase regulatory approval process adopted in those specific contexts. Further research on approval regulation under uncertainty could also shed light on the more general design of organizations to foster learning, with implications for broader principal-agent environments.

APPENDIX A

A.1 Omitted proofs

Proof of Proposition 2. Using equation (1), we have

$$\frac{\partial W}{\partial n} = E[\theta] - \int_{\underline{\theta}}^{-k/n} \theta g(\theta)d\theta = \frac{W[k,n,g(\theta)]}{n} - \frac{E[\theta]}{n} + \int_{\underline{\theta}}^{-k/n} \frac{k}{n} g(\theta)d\theta.$$

Evaluating this derivative at $k = \hat{k}$ and using $W[\hat{k}, n, g(\theta)] = 0$, we find that

$$\frac{\partial}{W}|\partial n_{k=\hat{k}} = -\frac{E[\theta]}{n} + \frac{\hat{k}}{n} G\left(-\frac{\hat{k}}{n}\right) > 0$$

under the assumption that $E[\theta] < 0$. Given equation (2), we conclude by the implicit function theorem that an increase in n leads to an increase in $\hat{k}$.

To investigate the effect of changes in the distribution of θ on $\hat{k}$, rewrite the expected social welfare when not blocking A ex ante as

$$W[k,n,g(\theta)] = E[\theta] + \int_{\underline{\theta}}^{\overline{\theta}} \max\langle -k, n\theta\rangle g(\theta)d\theta.$$

It is then immediate to derive our last two comparative statics results. First, W increases with a first-order stochastic shift in the distribution of θ because θ and $\max\langle -k, n\theta\rangle$ are both increasing functions of θ. Second, from Rothschild and Stiglitz (1970) we know that a mean-preserving spread in θ increases the expected value of the convex function of θ in the second term, while leaving $E[\theta]$ unaffected. We conclude that a first-order stochastic shift in θ and a mean-preserving spread in θ lead to an increase in $\hat{k}$. *Q.E.D.*

Example for Proposition 3. We now illustrate the sufficient conditions given in Proposition 3 through a simple numerical example of a situation in which an activity is efficiently taken ex ante only if the regulator commits not to review it ex post. Let $\pi(\theta) = z - b\theta$ for $z, b > 0$ and assume that θ is uniformly distributed. In this case,

established by Regulation 17/62) based on mandatory notification of agreements among competitors with a new regime of ex post control, mostly implemented by competition authorities at the state level. See Loss et al. (2008) for more details.

$E[\theta] > 0$ if and only if $\underline{\theta} + \overline{\theta} > 0$. $E[\pi(\theta)] > 0$ if and only if $z > -b(\underline{\theta} + \overline{\theta})/2$. To ensure that $E[\pi] + n\int_{-k/n}^{\overline{\theta}} \pi(\theta)g(\theta)d\theta - kG(-k/n) < 0$, so that the firm does not want to file for A in the presence of ex post review, z cannot be too large. In this example, this condition is $z < \{b[\overline{\theta}^2 n(1+n) - \underline{\theta}^2 n - k^2] - 2k(k + \underline{\theta}n)\}/\{2n[\overline{\theta}(1+n) - \underline{\theta} + k]\}$. To show that it is possible to satisfy all of these conditions on z simultaneously, let $\underline{\theta} = -1/2, \overline{\theta} = 1, k = 1/2, b = 1/4, n = 2$. Thus, we need $z > 1/16$ for A to be profitable in expectation when the regulator commits not to reverse it ex post. For ex ante review to deter A, we need $z < 29/256$. Since $29/256 > 1/16$, all of the conditions are satisfied.[21]

Proof of Proposition 4. If $E[\pi] + n\Pr(\theta \geq -k/n)E[\pi|\theta \geq -k/n] < 0$, $E[\pi] > 0$ and $E[\theta] > 0$, then we know from Proposition 1 that the regulator wants the firm to propose A. However, the firm will never undertake activity A since it generates negative profits at the current cost. Thus, R needs to commit to some ex post review policy other than what would be ex post optimal. Approval for $\theta > \theta^c$ will induce the firm to propose A by definition of θ^c. To prove that this is the socially optimal way to meet the zero profit constraint, assume that the regulator chooses to approve the proposal in some region $[\theta_1, \theta_2]$ for $\theta_2 \leq -k/n$. Because approval in $[\theta_1, \theta_2]$ reduces social welfare, this must be done to help meet the zero profit constraint by providing some positive level of profit $\int_{\theta_1}^{\theta_2}[n\pi(\theta) + k]g(\theta)d\theta =: z > 0$, given that $\pi(-k/n) > 0$ and $\pi'(\theta) < 0$. We will show that the socially optimal way to generate z in profits for F is increase θ_1 as much as possible, that is, until $\theta_2 = -k/n$ and, consequently, $\theta_1 = \theta^c$.

To prove this, we begin by totally differentiating $\int_{\theta_1}^{\theta_2}[n\pi(\theta) + k]g(\theta)d\theta = z$ with respect to θ_1 and θ_2 to obtain $\frac{d\theta_2}{d\theta_1} = \frac{n\pi(\theta_1)+k}{n\pi(\theta_2)+k}\frac{g(\theta_1)}{g(\theta_2)}$. The welfare effect of approval in this region is $\int_{\theta_1}^{\theta_2}[n\theta + k]g(\theta)d\theta$. To compute the change in welfare due to a small change in θ_1, we differentiate the latter expression with respect to θ_1. Using $\frac{d\theta_2}{d\theta_1} = \frac{n\pi(\theta_1)+k}{n\pi(\theta_2)+k}\frac{g(\theta_1)}{g(\theta_2)}$, we obtain $g(\theta_1)\left[\frac{n\pi(\theta_1)+k}{n\pi(\theta_2)+k}(n\theta_2 + k) - (n\theta_1 + k)\right]$. Using $n\pi(\theta_2) + k > 0$ (otherwise there would be no reason not to overturn A at θ_2 because doing so would reduce profit and social welfare), this expression has the sign of $[n\pi(\theta_1) + k](n\theta_2 + k) - [n\pi(\theta_2) + k](n\theta_1 + k)$. Note that $[n\pi(\theta_1) + k]\ (n\theta_2 + k) - [n\pi(\theta_2) + k](n\theta_1 + k) = n\{[n\pi(\theta_2) + k](\theta_2 - \theta_1) + [\pi(\theta_1) - \pi(\theta_2)](n\theta_2 + k)\}$, where $[n\pi(\theta_2) + k](\theta_2 - \theta_1) > 0$ and $n\theta_2 + k < 0$.

Next, by the intermediate value theorem, we have $[n\pi(\theta_2) + k](\theta_2 - \theta_1) + [\pi(\theta_1) - \pi(\theta_2)](n\theta_2 + k) = [n\pi(\theta_2) + k](\theta_2 - \theta_1) + \pi'^{*})(\theta_1 - \theta_2)(n\theta_2 + k)$ for some $\theta^* \in [\theta_1, \theta_2]$. This is positive if and only if $[n\pi(\theta_2) + k] - \pi'^{*})(n\theta_2 + k) > 0$. Clearly, $[n\pi(\theta_2) + k] - \pi'^{*})[n\theta_2 + k] > n[\pi(\theta_2) - \pi'^{*})\theta_2]$ by the assumption $\pi' < 0$. By the intermediate value theorem, we have $\pi(\theta_2) = \pi(0) + \theta_2\pi'^{**})$ where $\theta^{**} \in [\theta_2, 0]$.

21 Note that these conditions can be compatible with the willingness of the firm not to voluntarily undo the action (see Footnote 3 for a discussion of this possibility). To guarantee that F never undoes A voluntarily, we need that $z > b\overline{\theta} - k/n$. For the numerical value in the example, we need $z > 0$ to ensure the firm never voluntarily undoes A, a condition that is automatically satisfied by the parameters considered here.

Using this, we obtain $n[\pi(\theta_2) - \pi'^{*})\theta_2] = n\{\pi(0) + \theta_2[\pi'^{**}) - \pi'^{*})]\}$. Since $\theta_2 < 0$ and $\theta^{**} > \theta^{*}$, this is positive if $\pi(0) > 0$ and $\pi'' < 0$ for $\theta < 0$.

Thus, increasing θ_1 has a positive effect on social welfare. This shows that it is optimal for R to set $\theta_2 = -k/n$. Next, notice that R will set θ_1 so as to minimize the expected welfare loss, resulting from the activity for $\theta_1 \leq \theta_1' \leq -k/n$, subject to the zero profit constraint. Hence, given that the constraint as well as the objective function are decreasing in θ_1, it is indeed optimal to set $\theta_1 = \theta^c$. *Q.E.D.*

Proof of Lemma 1. Define $\mathcal{P}(\theta) := \pi(\hat{a}, \theta) + (n-1)\pi(\overline{a}(\theta), \theta)$ and note that $\mathcal{P}(\theta_1) = -k$ by the definition of θ_1. Using optimality of $\overline{a}(\cdot)$ and the fact that $\pi_2(a, \theta) > 0$, it follows immediately that $\mathcal{P}$ is increasing in θ. When $\theta > \theta_2$, we have $\overline{a}(\theta) < \hat{a}$ because $\overline{a}(\theta_2) = \hat{a}$ and $\overline{a}(\theta)$ is decreasing in θ by $\pi_{12} < 0$. Thus, in this case the regulator does not undo the activity A and the firm chooses $\overline{a}(\theta)$ in period 1. If instead $\theta < \theta_2$, it is straightforward to see that the firm finds it optimal to prevent R from undoing activity A if and only if $\mathcal{P}(\theta) > -k$, which is equivalent to $\theta > \theta_1$, given that $\mathcal{P}$ is increasing. In conclusion, for $\theta \in [\underline{\theta}, \theta_1]$, F chooses $\overline{a}(\theta)$ and R undoes activity A in period 2; for $\theta \in (\theta_1, \theta_2]$, F chooses $\hat{a}$ in period 1, R does not undo A, and F chooses $\overline{a}(\theta)$ after period 1; for $\theta \in (\theta_2, \overline{\theta}]$, F chooses $\overline{a}(\theta)$ in all periods and R does not undo the activity A after period 1. Clearly, if $\theta_1 > \theta_2$, in the first period the firm chooses the best action $\overline{a}(\theta)$ for all θ. *Q.E.D.*

Proof of Proposition 5. By Lemma 1, social welfare given $\hat{a}$ is

$$\int_{\underline{\theta}}^{\theta_1(\hat{a})} [\theta - k]g(\theta)d\theta + \int_{\theta_1(\hat{a})}^{\theta_2(\hat{a})} [\omega(\hat{a}, \theta) + n\theta]g(\theta)d\theta + \int_{\theta_2(\hat{a})}^{\overline{\theta}} (1+n)\theta g(\theta)d\theta.$$

Differentiating this with respect to $\hat{a}$ yields

$$\theta_1'(\hat{a})[\theta_1(\hat{a}) - \omega(\hat{a}, \theta_1(\hat{a})) - k - n\theta_1(\hat{a})]g(\theta_1(\hat{a})) + \int_{\theta_1(\hat{a})}^{\theta_2(\hat{a})} \omega_1(\hat{a}, \theta)g(\theta)d\theta \\ + \theta_2'(\hat{a})[\omega(\hat{a}, \theta_2(\hat{a})) - \theta_2(\hat{a})]g(\theta_2(\hat{a})).$$

The last term is zero since $\omega(\hat{a}, \theta_2(\hat{a})) = \theta_2(\hat{a})$ by the definition of θ_2 and ω. Thus, the net effect on welfare from a change in $\hat{a}$ is $\theta_1'(\hat{a})[\theta_1(\hat{a}) - \omega(\hat{a}, \theta_1(\hat{a})) - k - n\theta_1(\hat{a})]$ $g(\theta_1(\hat{a})) + \int_{\theta_1(\hat{a})}^{\theta_2(\hat{a})} \omega_1(\hat{a}, \theta)g(\theta)d\theta$. If we evaluate the welfare effect at $\hat{a}^*$, we obtain

$$\theta_1'(\hat{a}^*)[\theta_1(\hat{a}^*) - \omega(\hat{a}^*, \theta_1(\hat{a}^*))]g(\theta_1(\hat{a}^*)) + \int_{\theta_1(\hat{a}^*)}^{\theta_2(\hat{a}^*)} \omega_1(\hat{a}^*, \theta)g(\theta)d\theta, \qquad \text{(A1)}$$

using $-k - n\theta_1(\hat{a}^*) = 0$ by the definition of $\hat{a}^*$. To determine the sign of the first term, we first determine the sign of $\theta_1'(\hat{a})$. Total differentiation of $\pi(\hat{a}, \theta_1) +$ $(n-1)\tilde{\pi}(\theta_1) = -k$ with respect to θ_1 and $\hat{a}$ gives $\frac{d\theta_1}{d\hat{a}} = \frac{-\pi_1(\hat{a}, \theta_1)}{\pi_2(\hat{a}, \theta_1) + (n-1)\tilde{\pi}'(\theta_1)}$. The denominator is positive since $\pi_2(a, \theta) > 0$. The numerator is negative since $\hat{a}^* < \overline{a}(\theta_1)$ and $\pi_1(a, \theta) > 0$ for $a < \overline{a}(\theta)$. Thus, the first term of equation (A1) is positive since $\theta_1'(\hat{a}^*) < 0$ and $\theta_1(\hat{a}^*) - \omega(\hat{a}^*, \theta_1(\hat{a}^*)) < 0$ by the assumption that $\omega_1(a, \theta) < 0$. The second term is also negative by $\omega_1(a, \theta) < 0$. The overall effect is negative if and only if $g(\theta_1(\hat{a}^*)) < \frac{\int_{\theta_1(\hat{a}^*)}^{\theta_2(\hat{a}^*)} \omega_1(\hat{a}^*,\theta)g(\theta)d\theta}{\omega(\hat{a}^*,\theta_1(\hat{a}^*))-\theta_1(\hat{a}^*)} \frac{\pi_2(\hat{a}^*,\theta_1(\hat{a}^*))+(n-1)\tilde{\pi}'(\theta_1(\hat{a}^*))}{\pi_1(\hat{a}^*,\theta_1(\hat{a}^*))}$, as claimed. *Q.E.D.*

Proof of Proposition 6. The expected social welfare is

$$\breve{W}\ [k, n, g(\theta)] := (1+n)E[\theta] + \pi(k,\theta) - \pi(\bar{k},\theta) + \int_{\underline{\theta}}^{-k/n} [-n\theta - k]g(\theta)d\theta. \tag{A2}$$

Thus, the socially optimal level of k is given implicitly by

$$\frac{\partial \breve{W}}{\partial k} = \pi_1(k^*,\theta) - G\left(-\frac{k^*}{n}\right) = 0. \tag{A3}$$

Assuming that $\pi_1(0,\theta) - G(0) > 0$, we have an interior solution $0 < k^* < \bar{k}$. As far as the firm is concerned, the modified profit function is

$$\breve{\Pi}\ [k, n, g(\theta)] := \pi(k,\theta) + nE[\pi] + \int_{\underline{\theta}}^{-k/n} [-n\pi(\theta) - k]g(\theta)d\theta, \tag{A4}$$

recalling that R cannot commit ex ante. It thus follows that the optimal level of k for the firm solves

$$\frac{\partial \Pi}{\partial k} = \left[\pi\left(-\frac{k^{**}}{n}\right) + \frac{k^{**}}{n}\right] g\left(-\frac{k^{**}}{n}\right) + \pi_1(k^{**},\theta) - G\left(-\frac{k^{**}}{n}\right) = 0. \tag{A5}$$

Under Condition 1, comparing equations (A3) and (A5) we conclude that $k^{**} \geq k^*$. *Q.E.D.*

A.2 Multiple equilibria for the discipline effect

If the cutoff level of the action for reversal, $\hat{a}$, is set above $\hat{a}^*$, then both θ_1 and θ_2 shift down because they are both decreasing in a. Moreover, as $\hat{a}^*$ is set such that $\theta_1 = -k/n$, setting $\hat{a}$ above $\hat{a}^*$ implies that $\theta_1 < -k/n$. In such an equilibrium, if R observes $\hat{a}$ in period 1, then it will believe that $\theta \in (\theta_1, \theta_2]$ where $\theta_1 < -k/n < \theta_2$. By allowing A to continue in this region, R risks approving a welfare-reducing action but approves an action that is welfare-increasing (taking into account the reversal cost). If θ_1 is close enough to $-k/n$, because $\hat{a}$ is close enough to $\hat{a}^*$, then the probability that allowing the activity to continue is welfare-reducing will be quite low. Thus, R will still find it ex post optimal not to undo A. The next proposition describes the possible equilibria in this signaling game. Define first the functions $\theta_1(\hat{a})$ and $\theta_2(\hat{a})$ as in Lemma 1, that is, such that $\tilde{\pi}(\theta_1) - k = \pi(\hat{a},\theta_1) + n\tilde{\pi}(\theta_1)$ and $\bar{a}(\theta_2) = \hat{a}$. Note that both $\theta_1(a)$ and $\theta_2(a)$ are decreasing in a by the assumption that $\pi_{12} < 0$.

Proposition 7. Let $\hat{a}'$ be implicitly defined by $\frac{n}{G(\theta_2(\hat{a}'))-G(\theta_1(\hat{a}'))}\int_{\theta_1(\hat{a}')}^{\theta_2(\hat{a}')} \theta g(\theta)d\theta = -k$. For any $\hat{a} \in [\hat{a}^*, \hat{a}']$, if $\pi_2(a,\theta) > 0$ and $n > 1$, then the regulator undoes activity A in period 1 if and only if $a > \hat{a}$ and the firm chooses $\hat{a}$ for $\theta \in (\theta_1(\hat{a}), \theta_2(\hat{a})]$ and chooses $\overline{a}(\theta)$ for $\theta \in [\underline{\theta}, \theta_1(\hat{a})]$ and $\theta \in (\theta_2(\hat{a}), \overline{\theta}]$. The regulator believes that $\theta \in [\underline{\theta}, \theta_1(\hat{a})]$ if $a > \hat{a}$.

This proposition indicates that there is a continuum of equilibria when the regulator cannot commit to a policy of when to undo activity *A*. One of these equilibria, the one in which $\hat{a} = \hat{a}^*$, maximizes ex post social welfare.

Proof of Proposition 7. If *R* undoes the activity if and only if $a > \hat{a}$, then Lemma 1 establishes *F* will act as stated in the Proposition. If *F* does play this strategy, then *R*'s payoff (which is simply social welfare) after seeing $a = \hat{a}$ is given by $\frac{n}{G(\theta_2(\hat{a}))-G(\theta_1(\hat{a}))}\int_{\theta_1(\hat{a})}^{\theta_2(\hat{a})} \theta g(\theta)d\theta$ if it does not undo activity *A* and by—*k* if it does undo activity *A*. Since $\hat{a} < \hat{a}'$, and $\theta_1(a)$ and $\theta_2(a)$ are decreasing in *a*, we have

$$\frac{n}{G(\theta_2(\hat{a})) - G(\theta_1(\hat{a}))}\int_{\theta_1(\hat{a})}^{\theta_2(\hat{a})} \theta g(\theta)d\theta \geq \frac{n}{G(\theta_2(\hat{a}')) - G(\theta_1(\hat{a}'))}\int_{\theta_1(\hat{a}')}^{\theta_2(\hat{a}')} \theta g(\theta)d\theta = -k.$$

Therefore, *R*'s payoff is weakly greater if it allows *A*. When observing $a < \hat{a}$, *R* infers $\theta \in (\theta_2(\hat{a}), \overline{\theta}]$, generating a larger payoff from allowing *A*. When observing $a > \hat{a}$, *R* believes $\theta \in [\underline{\theta}, \theta_1(\hat{a})]$, ensuring a larger payoff from undoing *A*. *R*'s beliefs follow from *F*'s strategy except for $a \in (\hat{a}, \overline{a}(\theta_1(\hat{a})))$. Note that believing that $\theta \in [\underline{\theta}, \theta_1(\hat{a})]$ when seeing this off-equilibrium action does not violate the intuitive criterion (Cho and Kreps 1987). *Q.E.D.*

ACKNOWLEDGEMENTS

This is a slight revision of the manuscript that won the 2009 Robert F. Lanzillotti Prize for the best paper in antitrust economics accepted for presentation at the International Industrial Organization Conference. An earlier draft of the article was circulated in 2008 under the title "Policy Timing under Uncertainty: Ex Ante versus Ex Post Merger Control." We are grateful for comments by our discussants (Louis Kaplow, Alvin Klevorick, Kai-Uwe Kühn, and Rune Stenbacka), as well as by Amy Ando, Andrei Barbos, Umberto Garfagnini, Jerry Hausman, Emeric Henry, Lynne Kiesling, Niko Matouschek, Volker Nocke, Martin Peitz, Patrick Rey, Scott Stern, Andrew Thomson, Helen Weeds, and seminar participants at Duke, HEC Montreal, the NBER Law and Economics Program, the Searle Conference on Antitrust Economics and Policy at Northwestern University, the NIE Conference at the UK Competition Commission in London, the 2009 International Industrial Organization Conference in Boston, and the ZEW Conference on Ex-Post Evaluation of Competition Policy in Frankfurt. We thank the coeditor and referees for their constructive comments and Federico Pessina for excellent research assistance in updating the draft.

FUNDING

Financial support from the Searle Center on Law, Regulation, and Economic Growth is gratefully acknowledged.

CONFLICT OF INTEREST

Wickelgren occasionally consults on antitrust cases.

REFERENCES

Alonso, Ricardo, and Niko Matouschek. 2008. "Optimal Delegation," 75 *Review of Economic Studies* 259–93.

Amador, Manuel, George-Marios Angeletos, and Ivan Werning. 2006. "Commitment vs. Flexibility," 74 *Econometrica* 365–96.

Argentesi, Elena, Paolo Buccirossi, Emilio Calvano, Tomaso Duso, Alessia Marrazzo, and Salvatore Nava. 2021. "Merger Policy in Digital Markets: An Ex-Post Assessment," 17 *Journal of Competition Law and Economics* 95–140.

Armstrong, Mark, and John Vickers. 2010. "A Model of Delegated Project Choice," 78 *Econometrica* 213–44.

Arrow, Kenneth J., and Anthony C. Fisher. 1974. "Environmental Preservation, Uncertainty, and Irreversibility," 88 *Quarterly Journal of Economics* 312–19.

Barros, Pedro Pita. 2003. "Looking behind the Curtain—Effects from Modernization of European Union Competition Policy," 47 *European Economic Review* 613–24.

Besanko, David, and Daniel Spulber. 1993. "Contested Mergers and Equilibrium Antitrust Policy," 9 *Journal of Law, Economics and Organization* 1–29.

Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation," 67 *Econometrica* 349–74.

Bose, Arup, Bebashis Pal, and David E. Sappington. 2016. "On the Merits of Antitrust Liability in Regulated Industries," 59 *Journal of Law and Economics* 359–92.

Cho, In-Koo, and David M. Kreps. 1987. "Signaling Games and Stable Equilibria," 102 *Quarterly Journal of Economics* 179–221.

European Commission, Case No COMP/M.7217 - Facebook/ WhatsApp: Commission decision pursuant to Article 6(1)(b) of Council Regulation No 139/2004 (2014).

Daripa, Arup, and Simone Varotto. 2010. "Ex-Ante versus Ex-Post Regulation of Bank Capital," in Lloyd P. Blenman, Harold A. Black, and Edward J. Kane, eds., *Banking and Capital Markets: New International Perspectives*. Singapore: World Scientific.

Farrell, Joseph, and Michael L. Katz. 2006. "The Economics of Welfare Standards in Antitrust," 2 *Competition Policy International* 3–28.

Henry, Claude. 1974. "Investment Decisions under Uncertainty: The Irreversibility Effect," 64 *American Economic Review* 1006–12.

Henry, Emeric, and Marco Ottaviani. 2019. "Research and Approval Process: The Organization of Persuasion," 109 *American Economic Review* 911–55.

Henry, Emeric, Marco Loseto, and Marco Ottaviani. 2022. "Regulation with Experimentation: Ex Ante Approval, Ex Post Withdrawal, and Liability," 68 *Management Science* 5330–47.

Holmström, Bengt. 1984. "On the Theory of Delegation," in Marcel Boyer and Richard Kihlstrom, eds., *Bayesian Models in Economic Theory*. Amsterdam: Elsevier.

Katz, Michael L., and Howard A. Shelanski. 2007. "Merger Analysis and the Treatment of Uncertainty: Should We Expect Better?" 74 *Antitrust Law Journal* 537–74.

Kolstad, Charles D., Thomas S. Ulen, and Gary V. Johnson. 1990. "Ex Post Liability for Harm vs. Ex Ante Safety Regulation: Substitutes or Complements?" 80 *American Economic Review* 888–901.

Kotowski, Maciej H., David A. Weisbach, and Richard J. Zeckhauser. 2014. "Rules and Standards When Compliance Costs Are Private Information," 43 *Journal of Legal Studies* S297–29.

Kwoka, John, and Tommaso Valletti. 2021. "Unscrambling the Eggs: Breaking Up Consummated Mergers and Dominant Firms," 30 *Industrial and Corporate Change* 1286–1306.

Loss, Frédéric, Estelle Malavolti-Grimal, Thibaud Vergé, and Fabian Bergès-Sennou. 2008. "European Competition Policy Modernization: From Notifications to Legal Exception," 52 *European Economic Review* 77–98.

Mason, Robin, and Helen Weeds. 2013. "Merger Policy, Entry, and Entrepreneurship," 57 *European Economic Review* 23–38.

Ottaviani, Marco, and Abraham L. Wickelgren. 2011. "Ex Ante or Ex Post Competition Policy? A Progress Report," 29 *International Journal of Industrial Organization* 356–59.

Patel, Menesh S. 2020. "Merger Breakups," 5 *Wisconsin Law Review* 975–1042.

Rey, Patrick. 2003. "Towards a Theory of Competition Policy," in Mathias Dewatripont, Lars Peter Hansen, and Steven J. Turnovsky, eds., *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*. Cambridge: Cambridge University Press.

Rothschild, Michael, and Joseph E. Stiglitz. 1970. "Increasing Risk: I. A Definition," 2 *Journal of Economic Theory* 225–43.

Salop, Steven C. 2016. "Modifying Merger Consent Decrees to Improve Merger Enforcement Policy," 31 *Antitrust* 15–20.

Schwartzstein, Joshua, and Andrei Shleifer. 2013. "An Activity-Generating Theory of Regulation," 56 *Journal of Law and Economics* 1–38.

Shavell, Steven. 1984. "A Model of the Optimal Use of Liability and Safety Regulation," 5 *RAND Journal of Economics* 271–80.

Shubik, Martin, and Richard Levitan. 1980. *Market Structure and Behavior*. Cambridge, MA: Harvard University Press.

Strulovici, Bruno. 2010. "Learning While Voting: Determinants of Collective Experimentation," 78 *Econometrica* 933–71.

Taylor, Curtis, and Huseyin Yildirim. 2011. "Subjective Performance and the Value of Blind Evaluation," 78 *Review of Economic Studies* 762–94.

Viscusi, W. Kip, John M. Vernon, and Joseph E. Harrington. 1995. *Economics of Regulation and Antitrust*. Cambridge, MA: MIT Press.

Wald, Abraham. 1945. "Sequential Tests of Statistical Hypotheses," 16 *Annals of Mathematical Statistics* 117–86.

Weeds, Helen. 2002. "Strategic Delay in a Real Options Model of R&D Competition," 69 *Review of Economic Studies* 729–47.