

NBER WORKING PAPER SERIES

PROPOSED MERGERS WHERE EFFICIENCIES ARE NEEDED MOST MIGHT
BE THE LEAST LIKELY TO DELIVER THEM

Robert D. Metcalfe
Alexandre B. Sollaci
Chad Syverson

Working Paper 35221
<http://www.nber.org/papers/w35221>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2026

We thank Emir Kamenica for helpful conversations. Syverson discloses that he consults as an expert in antitrust matters. The views expressed in this paper do not necessarily represent those of the IMF, its Executive Board, IMF management, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Robert D. Metcalfe, Alexandre B. Sollaci, and Chad Syverson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Proposed Mergers Where Efficiencies Are Needed Most Might Be the Least Likely to Deliver Them

Robert D. Metcalfe, Alexandre B. Sollaci, and Chad Syverson

NBER Working Paper No. 35221

May 2026

JEL No. L0

ABSTRACT

Mergers are commonly evaluated by weighing their expected market power effects against any efficiency gains they create. The larger the market power effect of a proposed merger, the larger must be any efficiencies for it to raise social welfare. We show selection into merger proposal distorts the observed relationship between market power and efficiency effects. Even if market power and efficiency gains are independent (or even positively correlated) across all potential mergers, they will generally be negatively related among proposed mergers. This is because parties propose to merge only if the merger's expected profitability exceeds a threshold, so the underlying components of profitability become substitutes in clearing that hurdle. It does not rely on managerial bias, behavioral frictions, or strategic misrepresentation. We demonstrate this negative correlation is present under very general conditions when the two effects are uncorrelated among all mergers. We also characterize conditions where this still holds even in the presence of positive underlying correlations and firms' uncertainty about their own merger's profitability. Policies that might raise the selection hurdle for proposed mergers do not alleviate the negative correlation; indeed, they would exacerbate it. Our analysis has direct implications for interpreting empirical merger retrospectives and for evaluating efficiency defenses in antitrust policy.

Robert D. Metcalfe
Columbia University
School for International
and Public Affairs
and NBER
rdm2176@columbia.edu

Chad Syverson
The University of Chicago
Booth School of Business
and NBER
chad.syverson@chicagobooth.edu

Alexandre B. Sollaci
International Monetary Fund
abalduinosollaci@imf.org

1 Introduction

Mergers can have both market power and cost efficiency effects. Welfare-maximizing merger policy considers this tradeoff and seeks to prevent mergers where market power dominates while allowing those where cost efficiencies are bigger ([U.S. Department of Justice and Federal Trade Commission, 2023](#)). The larger the expected market power increase from a merger, the larger must be any cost efficiencies for it to be justified from a social welfare standpoint ([Williamson, 1968](#); [Farrell and Shapiro, 1990](#)).

Policymaking under this tradeoff is complicated by the fact that neither effect’s magnitude is observed with certainty at the time the merger is proposed, when legal challenges are generally made. A large literature has sought to develop and apply tools to measure mergers’ effects to aid in this discernment, but the practical reality is that residual uncertainty always remains.¹ The policy issue is how to view this decision under uncertainty.

This paper argues that evaluating a proposed merger’s social welfare effects is inherently complicated by the process through which potentially merging parties decide to consummate a merger. There is of course a fundamental misalignment: firms select mergers based on expected profitability, while policymakers evaluate them based on welfare. That is well known. Our point here is this disconnect interacts with the merger selection process to result in mergers with larger expected market power effects—where we would need to expect larger efficiency effects to justify them on welfare grounds—instead having *smaller* expected efficiency effects, all else equal. This negative correlation will exist even if market power effects and efficiency gains are unrelated (or, in some cases, even if they are positively related) across all potential mergers.

This perverse covariance between the two social welfare effects of proposed mergers is not the result of any managerial bias, overconfidence, or strategic subterfuge by merging parties. It is instead a straightforward, almost mechanical consequence of how mergers are agreed to by the merging parties.

Mergers are only proposed for consummation when the parties expect sufficient profitability net of any merger costs. Because profitability is an increasing function of both the market power and efficiency effects of a merger, and only their combination determines whether the profitability hurdle is met, this mathematically imparts a negative slope between the two effects on the profitability threshold. The truncation of potential mergers at this threshold can then in turn create a negative correlation between the two effects of *proposed* mergers.²

¹Just a few recent examples of work measuring mergers’ market power or efficiency effects include [Braginsky et al. \(2015\)](#); [Blonigen and Pierce \(2016\)](#); [Stiebale and Vencappa \(2018\)](#); [Cooper et al. \(2019\)](#); [Hortaçsu et al. \(2019\)](#); [Miller et al. \(2021\)](#); [Panhans and Taragin \(2023\)](#); [Chiang et al. \(2025\)](#); [Demirer and Karaduman \(2024\)](#).

²The selection mechanism we elucidate here is conceptually related to econometric models of selection and

Intuitively, if a proposed merger’s market power effect on the post-merger entity’s profits is small, the merger’s efficiency effect on profits must be especially large. Otherwise, the merger would not have been proposed. Conversely, if the market power effect is large, we should expect the merging parties to tolerate smaller efficiency effects and still want to merge, as market power gains would be enough to make the merger profitable.

This selection logic matters for evaluating mergers under uncertainty because proposed mergers that plausibly generate large market power gains will mechanically be associated with smaller expected efficiency gains. This clarifies why “efficiency defenses” may be less compelling on average precisely for the mergers that raise the largest market power concerns.

We proceed as follows. Section 2 contains a simple example of the selection mechanism. Section 3 then provides a general model. Next, Section 4 extends the logic to instances where the merging parties themselves do not know with certainty the market power and efficiency effects of their merger when they have to decide whether to merge. Section 5 discusses further policy implications and concludes.

2 A Simple Example of the Selection Mechanism

Suppose the profitability of a potential merger is increasing in two latent components, market power M and efficiencies (mnemonically, cost savings) C . Assume that each component is independently distributed according to a standard normal distribution and that profitability is additive:

$$\Pi = M + C.$$

A merger is proposed if, and only if, it is sufficiently profitable; i.e., it must clear some hurdle K :

$$\mathcal{P} \equiv \mathbb{1}\{\text{Propose Merger}\} = \mathbb{1}\{M + C \geq K\}.$$

Selection Result: Cost Savings and Market Power Are Negatively Related Fix the market power effect of a potential merger at $M = m$. Because C is independent of M , its marginal distribution remains $N(0, 1)$. For the merger to be proposed, the cost savings must be sufficiently large:

$$C \geq K - m.$$

truncation where conditioning on an index exceeding a threshold can induce correlations among components of that index even when none exist in the population (Roy, 1951; Heckman, 1979). This selection concept is also related to a similar selection issue in estimating worker or manager fixed effects in AKM type models (Metcalf et al., 2023).

Thus the expected cost savings of a proposed merger with market power effect m are

$$l(m) \equiv \mathbb{E}[C \mid M = m, \mathcal{P}] = \mathbb{E}[C \mid C \geq K - m].$$

This is the mean of a standard normal distribution truncated from below (the inverse Mills ratio):

$$l(m) = \frac{\varphi(K - m)}{1 - \Phi(K - m)}.$$

The derivative of the truncated-from-below inverse Mills ratio with respect to the truncation point is always positive. Because the truncation point here, $K - m$, is decreasing in m , the derivative of $l(m)$ must be negative:

$$l'(m) = \frac{\varphi(K - m)}{1 - \Phi(K - m)} \left(\frac{\varphi(K - m)}{1 - \Phi(K - m)} - (K - m) \right) < 0$$

Thus the expected efficiency effect of a proposed merger falls as the merger's market power effect rises. Perversely from an antitrust policy standpoint, proposed mergers that will tend to raise market power more will have smaller expected efficiency effects. This is true even though the sizes of the two effects are independently distributed among all potential mergers. Selection on profitability and the fact that both market power and efficiency effects raise profitability create their negative correlation among proposed mergers.

2.1 The Selection Result in One Picture

The above result is easily seen graphically. Figure 1 plots the market power and efficiency effects for a set of potential mergers. The two effects are uncorrelated, distributed joint standard normal with correlation parameter $\rho = 0$.

These potential mergers are evaluated by merger partners who decide to move forward if the merger would be sufficiently profitable, i.e., if $M + C \geq K$ (here, $K = 1$). Profitable mergers are those with M, C pairs above and to the right of the downward-sloping profitability threshold; i.e., the orange points.

It is visually apparent that, even though market power and efficiency effects are uncorrelated overall, among mergers actually proposed, there is a negative correlation between the two effects. This is because the profitability threshold is downward-sloping. Proposed mergers with large market power effects tend to have small efficiency effects and still be profitable.

Figure 1: Selection into Mergers When Merger Effects are Initially Uncorrelated




Figure shows the distribution of potential mergers' M and C when both are i.i.d. $\mathcal{N}(0, 1)$. Also shown is the profitability threshold $M + C \geq K = 1$. Potential mergers above and to the right of this threshold will be proposed. The conditional mean of C given M among selected mergers traces a downward-sloping function.

3 General Model

The mechanism does not rely on linearity of profits or normality of merger characteristics. (Though we show further below there are natural reasons why profits might be linear in market power and efficiencies.) We discuss here how the same selection-based negative relationship between proposed mergers' expected market power and efficiency effects exists under more general conditions.

Assumption 1. (*Monotone profits.*) Each potential merger is characterized by (M, C) drawn from a joint distribution with a density on $\mathbb{R}^2$. The profitability of a proposed merger is given by function $\Pi = \pi(M, C)$, which is continuously differentiable and strictly increasing in both arguments:

$$\pi_M(m, c) > 0, \quad \pi_C(m, c) > 0 \quad \forall (m, c).$$

As before, a merger is proposed if it is sufficiently profitable:

$$\mathcal{P} \equiv \mathbb{1}\{\pi(M, C) \geq K\}.$$

Proposition 1. *The Profitability Frontier for Proposed Mergers is Downward-Sloping*

Proof. Fix $M = m$. Because π is strictly increasing in C , there exists a unique cutoff $h(m)$ such that

$$\pi(m, h(m)) = K \tag{1}$$

and $\pi(m, C) \geq K$ for any $C \geq h(m)$. Differentiate implicitly:

$$\pi_M(m, h(m)) + \pi_C(m, h(m))h'(m) = 0 \quad \Rightarrow \quad h'(m) = -\frac{\pi_M(m, h(m))}{\pi_C(m, h(m))} < 0.$$

□

The merger proposal boundary is still downward sloping under this general profit function. The result comes from the fact that merger profitability is a positive function of both market power and efficiency effects, so a larger market power effect reduces the minimum required efficiency effect to clear the hurdle.

Next, we discuss sufficient conditions for the negative relationship between cost savings and profitability in proposed mergers to hold. We start by showing that this result is true under Assumption 1 and if M and C are independent. We then relax the independence assumption and derive sufficient conditions for the negative relationship to hold when M and C are jointly normal.

3.1 Sufficient Conditions for Negative Selection

Independence between M and C

Let us examine how this negative proposal boundary affects the correlation of market power and efficiency effects among proposed mergers. We first assume M and C are independent among all potential mergers and denote the density of C as f_C . Conditional on $M = m$, the proposal process truncates C from below at the cutoff $h(m)$, defined by equation (1). The conditional mean efficiency effect among proposed mergers is

$$l(m) \equiv \mathbb{E}[C \mid M = m, \mathcal{P}] = \frac{\int_{h(m)}^{\infty} c f_C(c) dc}{\int_{h(m)}^{\infty} f_C(c) dc} \equiv \frac{A(m)}{B(m)},$$

where $A(m)$ is the truncated first moment and $B(m)$ the survival probability. Using Leibniz's rule, we have

$$A'(m) = -h'(m) h(m) f_C(h(m)), \quad \text{and} \quad B'(m) = -h'(m) f_C(h(m))$$

and differentiating $l(m)$ gives

$$l'(m) = \frac{A'(m)B(m) - A(m)B'(m)}{B(m)^2} = \frac{-h'(m)f_C(h(m))(h(m)B(m) - A(m))}{B(m)^2}.$$

Note that

$$A(m) - h(m)B(m) = \int_{h(m)}^{\infty} (c - h(m))f_C(c) dc > 0 \quad \Rightarrow \quad h(m)B(m) - A(m) < 0.$$

Since $h'(m) < 0$ (Proposition 1) and $f_C(h(m)) > 0$, it follows that

$$l'(m) < 0.$$

Thus, under Assumption 1 and independence of the market power and efficiency effects, the expected size of the efficiency effect among proposed mergers is decreasing in the market power effect.

Joint Normality of M and C

The results above demonstrate how, under very general conditions, selection yields a negative relationship between expected market power and efficiency effects among proposed mergers whenever the two effects are uncorrelated among all potential mergers. When market power and efficiency effects are instead correlated among all potential mergers, their relationship among proposed mergers depends on the properties of their specific distributions. It is not possible to derive general conditions that ensure a negative correlation among proposed mergers in this case. However, it is possible and instructive to do so for specific merger profit functions and distributional structures for M and C . We work through such a case here.

We start by introducing some structure to the joint profit of firms within a potential merger. Suppose firms A and B compete in a market, and each can produce y units of a given good with cost $C_i(y)$, for $i = A, B$. Assuming the cost function is differentiable, define firm i 's markup as

$$\mu_i(y) = \frac{p_i(y) - C'_i(y)}{C'_i(y)},$$

where p_i is the price charged by company i . The firm's profit can therefore be written as

$$\pi_i(y) = p_i y - C_i(y) = (1 + \mu_i) C'_i(y) y - C_i(y).$$

It is useful to define $R_i(y) \equiv C'_i(y) y$ as the revenue that firm i would get under perfect competition, when price equals marginal cost.

A merger between companies A and B would create a new firm, N , that has a potentially different markup and technology. Its profit function is given by

$$\pi_N(y) = (1 + \mu_N) R_N(y) - C_N(y) = \mu_N R_N + (R_N - C_N).$$

We also assume the merger between firms A and B incurs a cost that is proportional to the profits of both firms: $(e^K - 1) \times (\pi_A(y_A) + \pi_B(y_B))$, where the parameter K controls the size of the cost (note that when $K = 0$ there is no cost). This can be thought of as legal and regulatory hurdles to proposing a merger, the time and resources firms invest into due diligence and negotiating terms, and so on, all relative to the combined profits of the firms.

If y_i^* denotes the optimal production level for firm i , a merger is proposed if $\pi_N(y_N^*) - [\pi_A(y_A^*) + \pi_B(y_B^*)] \geq (e^K - 1)(\pi_A(y_A^*) + \pi_B(y_B^*))$ or (suppressing the y_i^* arguments),

$$\frac{\pi_N}{\pi_A + \pi_B} \geq e^K.$$

Lemma 1 shows that the conditions under which firms A and B would propose a merger can be approximated by a linear selection rule (the proof is found in Appendix A.1).

Lemma 1. *A first order approximation of the condition under which firms A and B find it profitable to propose a merger induces the selection rule*

$$\mathcal{P} \equiv \mathbb{1}\{\pi(M, C) \geq K\} \equiv \mathbb{1}\{M + C \geq K\}.$$

where

$$M = \alpha \log \left(\frac{\mu_N}{\mu_A + \mu_B} \right) \quad \text{and} \quad C = (1 - \alpha) \log \left(\frac{R_N - C_N}{(R_A - C_A) + (R_B - C_B)} \right),$$

and

$$\alpha = \frac{\mu_A R_A + \mu_B R_B}{(\mu_A R_A + \mu_B R_B) + [(R_A - C_A) + (R_B - C_B)]}.$$

Lemma 1 shows that $\pi(M, C)$ can be approximated by the sum of M and C and, more importantly, defines in intuitive terms what each of those components represents. $M \propto \log \left(\frac{\mu_N}{\mu_A + \mu_B} \right)$ is the increase in market power (measured as the size of markups) that firms

A and B expect to gain after the merger. Note that it is only positive when the markup of the merged firm exceeds the sum of markups of the two original companies.

In contrast, $C \propto \log \left(\frac{R_N - C_N}{(R_A - C_A) + (R_B - C_B)} \right)$ measures the gains in productivity and cost savings that the merged firm would enjoy relative to firms A and B . Recall that R_i is the revenue made by firm i under perfect competition, so a positive C indicates the merged firm is either has lower marginal cost than the two original companies or is able to obtain efficiencies through reducing fixed costs.

The factor α is the share of pre-merger joint profits attributable to markups. When firms A and B enjoy high market power, they will place a higher weight on the potential market power gains from the merger when deciding whether to propose it. On the other hand, if the two firms have low market power but high competitive revenue, they will place more emphasis on the merger's efficiency gains when deciding whether to propose it. This can be tied to the pre-existing market structure: in markets that are already highly competitive, firms might not be able to exploit much market power through mergers. However, if the two firms already enjoy a considerable amount of market power, their ability to maintain or increase that power after a merger is an important consideration.

Next, let us consider a particular distribution for the terms M, C :

Assumption 2. (*Joint normality.*) Each potential merger is characterized by a draw (M, C) from a joint normal distribution:

$$\begin{pmatrix} M \\ C \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho\sigma \\ \rho\sigma & \sigma^2 \end{pmatrix} \right).$$

Note: we can impose $\mathbb{E}[M] = \mathbb{E}[C] = 0$ and $\text{Var}(M) = 1$ without loss of generality when the merger selection rule is linear, as the hurdle for a proposed merger can always be rescaled to $k = \frac{K - \mathbb{E}[C] - \mathbb{E}[M]}{\text{Var}(M)}$.

At this stage, we should point out that while the joint normality of (M, C) is assumed in part because of analytical convenience, it is not unrealistic. In particular, it implies that ratios $\mu_N/(\mu_A + \mu_B)$ and $(R_N - C_N)/(R_A - C_A + R_B - C_B)$ are log-normally distributed: positive, right-skewed, and with heavy tails. This implies the market power and efficiency gains from mergers are low in most cases, but with significant heterogeneity across mergers and a nonnegligible probability of very large gains. These features are consistent with recent findings on the impacts of mergers (Blonigen and Pierce, 2016; Levonyan and Mengano, 2024).

Under these conditions, Proposition 2 derives the covariance between M and C in the set of proposed mergers under the two assumptions above (the proof can be found in Appendix

A.2).

Proposition 2. *If a merger is proposed when $\Pi = M + C \geq K$ and (M, C) are jointly normal, then*

$$\text{Cov}(M, C \mid \Pi \geq K) = \sigma\rho - \frac{(1 + \sigma\rho)(\sigma^2 + \sigma\rho)}{1 + \sigma^2 + 2\sigma\rho} \left[1 - \frac{\text{Var}(\Pi \mid \Pi \geq K)}{\text{Var}(\Pi)} \right].$$

The term in squared brackets can be written as

$$1 - \frac{\text{Var}(\Pi \mid \Pi \geq K)}{\text{Var}(\Pi)} = \lambda(\kappa)^2 - \kappa\lambda(\kappa)$$

where $\lambda(x) = \frac{\varphi(x)}{1 - \Phi(x)}$ is the inverse mills ratio and $\kappa = \frac{K}{\sqrt{1 + \sigma^2 + 2\sigma\rho}}$.

This proposition helps us understand two key properties of the selection mechanism for mergers. First, notice that the unconditional covariance between M and C is $\sigma\rho$, so the conditional covariance can be written as the sum of the unconditional covariance and a selection adjustment term. In addition, whenever $(1 + \sigma\rho)(\sigma^2 + \sigma\rho) \geq 0$, the effect of selection is to make the covariance more negative, because $1 + \sigma^2 + 2\sigma\rho = \text{Var}(\Pi) > 0$, and normality of Π guarantees that $\text{Var}(\Pi \mid \Pi \geq K) \leq \text{Var}(\Pi)$. This includes all instances where $\rho \geq 0$, in which case selection strictly reduces the covariance. For sufficiently small ρ (or sufficiently large K), the conditional covariance becomes negative *even if that correlation was positive among all potential mergers*. In most cases under normality, a negative correlation among all potential mergers will remain negative after the selection takes place.

Second, note that for any $\rho \in (-1, 1)$, $\sigma\rho < (1 + \sigma\rho)(\sigma^2 + \sigma\rho)/(1 + \sigma^2 + 2\sigma\rho)$. Furthermore, as K increases, the variance $\text{Var}(\Pi \mid \Pi \geq K)$ approaches zero. This means for any σ and ρ , there exists a value of K large enough to induce $\text{Cov}(M, C \mid \Pi \geq K) < 0$. In words, sufficiently high hurdles can make $\text{Corr}(M, C \mid \mathcal{P})$ negative among proposed mergers, even when $\rho > 0$ among all potential mergers. This result may seem counterintuitive at first, as a higher burden K selects mergers with high total profitability. However, because the selected set collapses toward the downward-sloping frontier, this induces a negative relationship between market power and efficiency.

This is a stark result: if an antitrust regulator, concerned about the natural negative correlation between market power and efficiency effects due to selection into merger, seeks to address this issue by raising the selection hurdle and make mergers less likely (here, raise K), it might serve only to exacerbate the negative correlation.

4 Selection When Profitability Signals Are Noisy

The above cases discuss situations where observers and policymakers may not know with certainty a merger's market power and efficiency effects, but merging parties do know. Here we consider instances where merging parties also receive noisy signals about the sizes of both effects. We show, assuming the two effects are uncorrelated among all potential mergers, selection on profitability still induces a negative relationship between the two effects among proposed mergers if the distribution of the two effects across all mergers is homoskedastic (i.e., every potential merger has the same-sized uncertainty about its market power and efficiency effects). If these effects instead have different variances across potential mergers, a negative covariance among proposed mergers is not guaranteed.

4.1 Homoskedastic Noisy Merger Effects

Let $\mathcal{I}$ denote the information available to the potentially merging parties, and define their posterior mean expectations of market power and efficiency effects as

$$U \equiv \mathbb{E}[M \mid \mathcal{I}], \quad V \equiv \mathbb{E}[C \mid \mathcal{I}].$$

Assumption 3. (*Homoskedastic Gaussian posteriors.*) Conditional on $\mathcal{I}$,

$$(M, C) \mid \mathcal{I} \sim \mathcal{N}((U, V), \Sigma),$$

where Σ is constant across candidate mergers.

The merging parties' expected profitability of a potential merger is

$$G(U, V) \equiv \mathbb{E}[\pi(M, C) \mid \mathcal{I}].$$

Under Assumption 3 and $\pi_M, \pi_C > 0$, G is strictly increasing in U and V :

$$\frac{\partial G}{\partial U} = \mathbb{E}[\pi_M(M, C) \mid \mathcal{I}] > 0, \quad \frac{\partial G}{\partial V} = \mathbb{E}[\pi_C(M, C) \mid \mathcal{I}] > 0.$$

Thus the proposal rule

$$\mathcal{P} = \mathbb{1}\{G(U, V) \geq K\}$$

defines a downward-sloping frontier $V \geq h(U)$. Under independence of (U, V) across potential mergers, the same truncation logic as in Section 3 implies $\mathbb{E}[V \mid U = u, \mathcal{P}]$ is decreasing in u . So again, even with the merging parties themselves facing uncertainty over the size

of the market power and efficiency effects, selection on profitability causes proposed mergers to have negatively correlated market power and efficiency effects. Formal details are in Appendix A.3.

4.2 Heteroskedastic Noisy Merger Effects

Things are more complicated when the posterior covariance of expected market power and efficiency effects varies across candidate mergers.

Assumption 4. (*Heteroskedastic Gaussian posteriors.*) Conditional on $\mathcal{I}$,

$$(M, C) \mid \mathcal{I} \sim \mathcal{N}((U, V), \Sigma),$$

where Σ varies across candidate mergers.

The parties' expectations of the profitability of a potential merger becomes

$$G(U, V, \Sigma) \equiv \mathbb{E}[\pi(M, C) \mid U, V, \Sigma].$$

For any *fixed* Σ , we have the homoskedastic result above: monotonicity of π implies $G(\cdot, \cdot, \Sigma)$ is increasing in U and V , generating a downward-sloping frontier $V \geq h_{\Sigma}(U)$.

However, unconditional patterns among proposed mergers can be altered by composition effects if the profit function is convex. This is because if Σ varies systematically with U in the proposed set of mergers, then $\mathbb{E}[V \mid U = u, \mathcal{P}]$ mixes different truncation frontiers. If π is locally convex, greater uncertainty can raise G via Jensen's inequality, changing which deals pass the hurdle and potentially attenuating or reversing the negative unconditional relationship.

Intuitively, with different amounts of uncertainty about market power and efficiency effects across potential mergers, there can be merger-specific hurdles. If the profit function $\pi(M, C)$ is convex, even locally, then more uncertainty over either effect raises the option value of merging. This results in higher-variance mergers clearing the profitability hurdle with lower posterior means.

Suppose higher expected market power U is associated with higher posterior variance Σ among proposed mergers (even if it isn't across all mergers). Then as the expected market power effect increases, the set of mergers that pass the hurdle are not just those with lower expected efficiency effects V as above. They also contain disproportionately more high-variance deals. So the expected efficiency effect mixes across different variance classes, each with its own truncation rule. This can confound the unconditional negative relationship.

Conversely, if merger profits are linear in the market power and efficiency effects (as in Lemma A.1), variance does not affect expected profits, so different Σ values do not change which mergers pass the hurdle; the frontier is the same for all. Another caveat is that if high variance deals tend to pass the hurdle under a convex profit function, once uncertainty is resolved, one might expect mergers with low realized market power/efficiency to exit the market. This can also induce a second form of selection on surviving mergers. Appendix A.4 states formal within- Σ results and discusses sufficient conditions (e.g., independence between Σ and (U, V)) under which the negative selection pattern survives aggregation.

5 Conclusions and Implications for Antitrust Policy

Selection into merger proposals induces substitutability between market power and efficiency components whenever both raise profitability. This selection mechanism can generate a negative relationship between market power and efficiency in observed mergers even if the components are independent or positively correlated in the population of potential mergers. The result is robust to nonlinear profitability and, under appropriate conditions, to noisy information environments.

These results have clear implications for evaluating prospective mergers' welfare effects. They also point out that observed market power and efficiency effects in retrospective merger analyses do not reflect the unconditional joint distributions of these effects, but rather a truncated distribution that is systematically shaped through a selection process. The marginal distributions of either effect are dependent on the level of the other.

The paper presents two additional results, albeit dependent on imposing more structure into the problem, with important implications for antitrust policy. First, we show that when two firms enjoy high market power, they will place a higher weight on the potential gains from market power when considering a merger. Conversely, firms whose profits depend less on markups will put a higher weight on the efficiency gains from mergers when considering one. This suggests that the welfare effects of proposed mergers may depend on the *current* market structure of firms that propose it, simply by selecting which firms choose to merge.

Second, we show that the correlation between market power and efficiency gains in proposed mergers can change depending on the size of merger costs, including regulatory costs. Antitrust regulations therefore act as a selection instrument over the joint distribution of merger characteristics. If antitrust agencies could reliably measure a proposed merger's market power and efficiency effects, this selection would be relatively inconsequential. The agency would simply allow mergers that would increase welfare and block the remainder. To the extent that uncertainty about the merger's effects remains, however, manipulating

regulatory costs to shape the selected set of proposed mergers could quite possibly backfire, inducing proposed mergers with relatively large market power effects relative to their efficiency gains.

References

- Blonigen, Bruce A. and Justin R. Pierce**, “Evidence for the Effects of Mergers on Market Power and Efficiency,” Working Paper 22750, National Bureau of Economic Research 2016.
- Braguinsky, Serguey, Atsushi Ohyama, Tetsuji Okazaki, and Chad Syverson**, “Acquisitions, Productivity, and Profitability: Evidence from the Japanese Cotton Spinning Industry,” *American Economic Review*, 2015, 105 (7), 2086–2119.
- Chiang, Hsin-Dao, Jordan Collison, Lorenzo Magnolfi, and Chad Sullivan**, “Enhancing the Merger Simulation Toolkit with Machine Learning,” *arXiv preprint arXiv:2506.05225*, 2025.
- Cooper, Zack, Stuart V. Craig, and Martin Gaynor**, “The Price Ain’t Right? Hospital Prices and Health Spending on the Privately Insured,” *Quarterly Journal of Economics*, 2019, 134 (1), 51–107.
- Demirer, Mert and Ömer Karaduman**, “Do Mergers and Acquisitions Improve Efficiency? Evidence from Power Plants,” Working Paper 32727, National Bureau of Economic Research 2024.
- Farrell, Joseph and Carl Shapiro**, “Horizontal Mergers: An Equilibrium Analysis,” *American Economic Review*, 1990, 80 (1), 107–126.
- Heckman, James J.**, “Sample Selection Bias as a Specification Error,” *Econometrica*, 1979, 47 (1), 153–161.
- Hortaçsu, Ali, Fernando Luco, Steven L. Puller, and Dongni Zhu**, “Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets,” *American Economic Review*, 2019, 109 (12), 4302–4342.
- Levonyan, Vardges and Paolo Mengano**, “Mergers and Market Power,” Working Paper 2024.
- Metcalf, Robert D, Alexandre B Sollaci, and Chad Syverson**, “Managers and productivity in retail,” Technical Report, National Bureau of Economic Research 2023.
- Miller, Nathan H., Gloria Sheu, and Matthew C. Weinberg**, “Oligopolistic Price Leadership and Mergers: The United States Beer Industry,” *American Economic Review*, 2021, 111 (10), 3123–3159.

- Panhans, Matthew T. and Charles Taragin**, “Consequences of Model Choice in Predicting Horizontal Merger Effects,” *International Journal of Industrial Organization*, 2023, *89*, 102985.
- Roy, Andrew D.**, “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 1951, *3* (2), 135–146.
- Stiebale, Joel and Dev Vencappa**, “Acquisitions, Markups, Efficiency, and Product Quality: Evidence from India,” *Journal of International Economics*, 2018, *112*, 70–87.
- U.S. Department of Justice and Federal Trade Commission**, “2023 Merger Guidelines,” <https://www.justice.gov/atr/2023-merger-guidelines> 2023. Issued December 18, 2023.
- Williamson, Oliver E.**, “Economies as an Antitrust Defense: The Welfare Tradeoffs,” *American Economic Review*, 1968, *58* (1), 18–36.

A Proofs and Discussions

A.1 Proof of Lemma 1

Rewrite each company's profits as the sum of their gains from market power and their competitive profit

$$\pi_i = \mu_i R_i + (R_i - C_i), \quad i = A, B, N$$

and define the proportional gains from merging

$$\theta_\mu = \frac{\mu_N R_N}{\mu_A R_A + \mu_B R_B} \quad \text{and} \quad \theta_R = \frac{R_N - C_N}{(R_A - C_A) + (R_B - C_B)}.$$

Note that θ_μ and θ_R are always weakly positive since markups and profits are evaluated at their optimal level. Furthermore $\theta_\mu > 1$ indicates that the merged company has more revenue from market power than the two original companies combined, indicating that the merger increased the degree of market power of the firm. Similarly, $\theta_R > 1$ indicates that the competitive profits in the merged company are larger than their combined analogs in the original firms, indicating either gains in productivity or cost savings (which could include fixed costs). It follows that

$$\frac{\pi_N}{\pi_A + \pi_B} = \alpha \theta_\mu + (1 - \alpha) \theta_R$$

where $\alpha = \frac{\mu_A R_A + \mu_B R_B}{(\mu_A R_A + \mu_B R_B) + [(R_A - C_A) + (R_B - C_B)]}$. Therefore

$$\begin{aligned} \log \left(\frac{\pi_N}{\pi_A + \pi_B} \right) &= \log (\alpha \theta_\mu + (1 - \alpha) \theta_R) \\ &\approx \alpha \log(\theta_\mu) + (1 - \alpha) \log(\theta_R) \end{aligned}$$

where the second equality follows from a first-order approximation of the logarithm.

Finally, take the logarithm of both sides in the merger proposal condition and plug in the result above to find

$$\underbrace{\alpha \log(\theta_\mu)}_M + \underbrace{(1 - \alpha) \log(\theta_R)}_C \geq K$$

□

A.2 Proof of Proposition 2

With a linear profit function, a merger is proposed iff $\Pi = M + C \geq K$, which truncates profits from below. Using the law of total covariance, the truncated covariance between M

and C can be written as

$$\text{Cov}(M, C \mid \Pi \geq K) = \text{Cov}[\mathbb{E}(M \mid \Pi), \mathbb{E}(C \mid \Pi) \mid \Pi \geq K] + \mathbb{E}[\text{Cov}(M, C \mid \Pi) \mid \Pi \geq K].$$

Since (M, C) are jointly normal, Π is also normally distributed and we have

$$\mathbb{E}[M \mid \Pi] = \mathbb{E}[M] + \frac{\text{Cov}(M, \Pi)}{\text{Var}(\Pi)}(\Pi - \mathbb{E}[\Pi]),$$

and similarly for $\mathbb{E}[C \mid \Pi]$. Recall that $\mathbb{E}[M] = \mathbb{E}[C] = 0$, which also implies $\mathbb{E}[\Pi] = 0$.

Next, and again using the properties of the multivariate normal distribution, the conditional covariance is

$$\text{Cov}(M, C \mid \Pi) = \text{Cov}(M, C) - \frac{\text{Cov}(M, \Pi)\text{Cov}(\Pi, C)}{\text{Var}(\Pi)}.$$

Plugging these into the expression above gives

$$\begin{aligned} \text{Cov}(M, C \mid \Pi \geq K) &= \frac{\text{Cov}(M, \Pi)\text{Cov}(C, \Pi)}{\text{Var}(\Pi)^2} \text{Var}(\Pi \mid \Pi \geq K) + \text{Cov}(M, C) - \frac{\text{Cov}(M, \Pi)\text{Cov}(\Pi, C)}{\text{Var}(\Pi)} \\ &= \text{Cov}(M, C) - \frac{\text{Cov}(M, \Pi)\text{Cov}(\Pi, C)}{\text{Var}(\Pi)} \left[1 - \frac{\text{Var}(\Pi \mid \Pi \geq K)}{\text{Var}(\Pi)} \right]. \end{aligned}$$

Lastly, the truncated variance under normality is

$$\text{Var}(\Pi \mid \Pi \geq K) = \text{Var}(\Pi) \left(1 + \frac{K}{\sqrt{\text{Var}(\Pi)}} \lambda \left(\frac{K}{\sqrt{\text{Var}(\Pi)}} \right) - \lambda \left(\frac{K}{\sqrt{\text{Var}(\Pi)}} \right)^2 \right),$$

where $\lambda(x) = \frac{\varphi(x)}{1-\Phi(x)}$, with φ and Φ representing the standard normal density and distribution, respectively. Define $\kappa = \frac{K}{\sqrt{\text{Var}(\Pi)}} = \frac{K}{\sqrt{1+\sigma^2+2\sigma\rho}}$. Combining the results above and plugging in the values from Assumption 2, we have

$$\text{Cov}(M, C \mid \Pi \geq K) = \sigma\rho - \frac{(1+\sigma\rho)(\sigma^2+\sigma\rho)}{1+\sigma^2+2\sigma\rho} [\lambda(\kappa)^2 - \kappa\lambda(\kappa)],$$

as desired. □

A.3 Homoskedistic Noisy Merger Effects

Under Assumption 3, the proposal set is an upper contour set of $G(U, V)$. Since G is strictly increasing in both arguments, for each u there is a cutoff $h(u)$ such that

$$\mathcal{P} = \{G(U, V) \geq K\} = \{V \geq h(U)\}, \quad h'(u) = -\frac{G_U}{G_V} < 0.$$

If U and V are independent (or weakly correlated) across candidate mergers, then conditional on $U = u$ selection truncates V from below at a cutoff decreasing in u , implying $\mathbb{E}[V \mid U = u, \mathcal{P}]$ is decreasing in u by the same argument as in Section 3.

A.4 Heteroskedastic Noisy Merger Effects

Under Assumption 4, for each fixed Σ the proposal condition defines $V \geq h_\Sigma(U)$ with $h'_\Sigma(u) < 0$. Therefore the negative selection logic holds *within each posterior-variance class* Σ (under independence/weak correlation of U and V conditional on Σ).

Unconditionally,

$$\mathbb{E}[V \mid U = u, \mathcal{P}] = \int \mathbb{E}[V \mid U = u, \mathcal{P}, \Sigma] dF(\Sigma \mid U = u, \mathcal{P}),$$

so the slope of the unconditional conditional-mean curve depends on how the mixing distribution $F(\Sigma \mid U = u, \mathcal{P})$ changes with u . A sufficient condition for inheriting the within- Σ negative relationship is that Σ be independent of U (or independent of (U, V)) and that selection does not induce strong dependence between Σ and U .