Merger Negotiations and Ex-Post Regret*

Dennis L. Gärtner [a] and Armin Schmutzler [b,†]

[a] Socioeconomic Institute, University of Zürich, Blümlisalpstr. 10,
8006 Zürich, Switzerland

[b] Socioeconomic Institute, University of Zürich, Blümlisalpstr. 10,
8006 Zürich, Switzerland

December 2008

*The authors thank David A. Miller, Benny Moldovanu, Andras Niedermayer, Andrea Patacconi, Andrew Postlewaite, Vasiliki Skreta and Dezsö Szalay, as well as seminar participants at the Econometric Society's 2005 World Congress, the 2005 Annual Meeting of the 'Verein für Socialpolitik', the Swiss IO Day 2006, and Tbilisi State University for very helpful discussions. The first author gratefully acknowledges financial support for this project from the Swiss National Science Foundation under Grant No. PBZH1-108484.

†Corresponding author (Tel.: +41 (0)44 634 2271, Fax: +41 (0)44 634 4907).
E-mail adresses: dennis.gaertner@soi.uzh.ch (D. L. Gärtner), arminsch@soi.uzh.ch (A. Schmutzler).

Abstract

We consider a setting in which two potential merger partners each possess private information pertaining both to the profitability of the merged entity and to stand-alone profits, and we investigate the extent to which this private information makes ex-post regret an unavoidable phenomenon in merger negotiations. To this end, we consider ex-post mechanisms, which use both players' reports to determine whether or not a merger will take place and what each player will earn in each case. When the outside option of at least one player is known, the efficient merger decision can be implemented by such a mechanism under plausible budget-balance requirements. When neither outside option is known, we show that the potential for regret-free implementation is much more limited, unless the budget balance condition is relaxed to permit money-burning in the case of false reports.

Keywords: Mergers, Mechanism Design, Asymmetric Information, Interdependent Valuations, Efficient Mechanisms.

JEL Classification: D82, L10, G34

1 Introduction

Mergers and acquisitions occur frequently in the corporate landscape. Nevertheless, it appears that often at least one party regrets the outcome of the transaction with the benefit of hindsight. Famous examples include the mergers of BMW and Rover, AOL and Time-Warner, or Mattel and The Learning Company, to name only a few. In many of these examples, the value of the new entity turns out to be lower than the combined value of its parts; in other cases, the owners of one of the firms lose out, as reflected in the well-known result that acquirers often overpay for the target (Andrade et al. [2]). Given the tremendous stakes in many merger decisions and the correspondingly large opportunity cost from accepting a merger at conditions that turn out to be unfavorable ex-post, it seems natural to ask whether such problems are inevitable.

A popular explanation for regret in merger transactions are *internal agency conflicts* as, for instance, self-interested managers engage in unprofitable mergers to expand their empire. More basically, regret in merger transactions may simply result from bad luck in situations of *symmetric uncertainty*: Mergers that are profitable in expectation may turn out to be unprofitable in the event of bad states. Symmetric uncertainty may include a wide range of different phenomena. For instance, the general economic downturn following September 11th 2001 has often been mentioned as a reason why some mergers that took place immediately before the attacks have not been successful. Less spectacularly, uncertainty about the future prospects of an industry can often lead to mistaken merger decisions.

However, there are clearly important cases of merger failure that cannot be traced back to symmetric uncertainty, but rather to *asymmetric information* between potential merger partners. For instance, it was only after the acquisition of CUC International, a marketer of discount membership clubs, that the investment banker Henry Silverman, owner of the hotel chain HFS, discovered that CUC had dramatically overstated its revenues. The revelation of this information after the merger led to a massive drop in the stock market value of the new firm Cendant (Business Week [8]). Similarly, in the merger between the German banks *Bayerische Vereinsbank* and *Hypobank*, the former was unaware that the partner's "balance sheet contained a time-bomb" until two years after the merger (The Economist [38]). Finally, environmental liabilities which may well be known to takeover targets lead to huge costs for the acquirer, for instance, in the energy industry (Brown et al. [6]). Beyond these drastic examples of partners concealing relevant information, quite generally parties often lack precise information on important characteristics of the potential partner, his productivity, corporate culture, etc.

Starting from these observations, our paper investigates the misalignment of interests between merging parties as a basic source of regret over merger decisions. To isolate this issue, we deliberately choose our assumptions so that neither internal agency conflicts nor symmetric uncertainty can play a role: We consider situations where ownership and control in each firm coincide, and where any information that is relevant to the merger decision is present with at least one of the two firms. However, we let parties possess different *pieces* of this information. We then ask: Is it possible to structure the merger negotiations and the decision process in a fashion such that no party regrets its behavior afterwards?

To illustrate our setting, consider the following two examples. First, suppose two firms are privately informed about their own (future) productivity. Typically, both a firm's stand-alone profits and the post-merger profits of an entity to which it belongs should be higher for more productive types. Moreover, if firms compete in the same market, each firm's productivity will usually have a negative impact on the competitor's stand-alone profits. As a second example, suppose two firms have been competing in a stable market environment over an extended period of time. Then, they may be expected to be reasonably well informed about those characteristics of the competitor that pertain to their stand-alone performance. However, they may have very limited knowledge about how they would fit together in the event of a merger, that is, how well the organizational cultures match. Thus, one might still expect substantial private information concerning post-merger profits, whereas—in contrast to the previous example—firms' outside options are common knowledge. This difference between the two examples (i.e., whether outside options are common knowledge or not) will turn out to be crucial to whether regret is avoidable.

In the setting described, merger negotiations have the function of assembling the privately held information, reaching a (preferably efficient) merger decision, and determining how profits are to be split among parties in the event of a merger. Using a mechanism-design approach, we formalize this negotiation process in terms of *merger mechanisms*, where agents (simultaneously) report their information (i.e., their 'type'), and each of the above decisions is made contingent on these reports.

Motivated by the question at hand, we restrict attention to merger mechanisms which result in no ex-post regret, that is, mechanisms such that truthful reporting is optimal for each party given any type of the other party and given truthful reporting on the other party's behalf. In addition, we impose natural budget-balance and individual-rationality conditions. By the former, agents' payoffs must sum to jointly realized profits for any outcome of the mechanism. By the latter, agents must receive at least their stand-alone profit for any equilibrium outcome of the mechanism.

We derive our first result for settings in which parties possess private information concerning the value of a merger, but where (at least) one party's stand-alone profits are commonly known. We show that a simple mechanism where, in the event of a merger, the party with commonly known outside options is bought out by the other firm achieves efficient implementation, that is, mergers take place if and only if they increase joint profits.

In the remainder of the paper, however, we show that the scope for extending this positive result to more general settings is extremely limited: For a large class of settings, not only is implementation of the *efficient* merger decision impossible, but only very trivial merger decision functions can avoid ex-post regret at all. Specifically, we derive this result for settings such that both players' stand-alone profits depend positively on own types, but non-positively on the competitor's type, whereas the payoffs of the merged firm depend positively on both types. This negative result then suggests that internal agency conflicts and symmetric uncertainty may not be the only reasons for regret over merger transactions: Even when merger parties possess all relevant information and there are no internal agency conflicts, it is often simply impossible to structure negotiations to avoid ex-post regret.

An important ingredient to this result is the aforementioned notion of budget balance, by which the mechanism must balance also *off equilibrium*, that is, if agents misreport their type. This property creates a problem of separating type profiles with identical merger profits, so that agents' equilibrium payoffs cannot differ between such type profiles. Intuitively, on the one hand, as payoffs can be conditioned on actual types only through total merger profits, agents cannot be punished selectively in the event that observed joint mergers are inconsistent with individual reports. On the other hand, budget balance prohibits collective punishment as it requires that merger profits are fully distributed to the merging parties.

This latter point leads us to ask whether the scope for implementation can be improved upon by allowing money to be burnt in the event of off-equilibrium reports. It turns out that, in this case, the efficient merger decision can always be reached in a regret-free and individually rational way. The practical relevance of this result hinges, however, on arguably strong assumptions concerning agents' ability to commit, or the existence of a third-party mediator such as a merchant bank.

From a mechanism-design perspective, the problem considered in this paper has several distinguishing features. First, parties' payments are conditionable on ex-post information: A merger mechanism must not only produce a merger decision, but also prescribe how joint profits are to be shared in the event of a merger. In the simplest case, this can be done by allocating fixed profit-shares in a merged entity. More sophisticated financial arrangements allow conditioning agents' payoffs on reported types and realized merger profits in an essentially arbitrary manner.[1] As has

[1] Practical examples include collars (Officer [33]), which use changes in stock prices to determine the partners' remuneration, and contingent value rights (Hietala et al. [17]), where sellers obtain put options on the shares of the new entity.

previously been noted, this 'contingent-payment' feature gives the mechanism designer additional degrees of freedom.[2] Second, the problem displays interdependent valuations in that agents' preferences over outcomes depend on each other's type. Third, requiring that merger negotiations always result in a budget-balanced outcome is natural, as there is no obvious candidate for a third-party residual claimant. Fourth and finally, as noted above, we require the mechanism to be regret-free.[3] Eventually, the combination of these features makes it impossible to directly relate our problem to the previous mechanism-design literature.[4]

This becomes compellingly transparent by contrasting our results to the classical problem of bilateral trade under asymmetric information. For instance, Myerson and Satterthwaite [32] show that, under fairly general conditions, there exists no ex-post efficient, individually rational Bayesian trading mechanism for indivisible goods. In contrast, it is a simple corollary of our positive result that once it is possible to condition payoffs on the realized value of trade—a very natural feature in a merger setting but perhaps unrealistic in most classical trade problems—then there in fact exists an efficient, budget-balanced, individually rational and regret-

[2]Hansen [16] (see also Crémer [10] and Riley [36]) show that, in an auction setting, conditioning payments on ex-post information allows the auctioneer to improve his expected revenues. Mezzetti [29] demonstrates in a more general setting that the use of ex-post information on payoffs can aid implementation of efficient decision rules.

[3]More generally, one may want to consider Bayesian-Nash implementation, that is, merger mechanisms such that truth-telling is optimal in *expected* terms, with expectations taken over the other party's possible types. Our restriction to regret-free mechanisms is immediately driven by the question at hand. Nonetheless, we should note that the recent literature provides several independent arguments for this restriction based on robustness concerns. Specifically, these concern the robustness of Bayesian mechanisms to players' having incorrect beliefs about types and other players' beliefs (see Bergemann and Morris [4], Chung and Ely [9], and the survey in Jehiel et al. [25]), as well as concerns about the necessity of *simultaneous* information revelation in Bayesian Mechanisms and the related problem of espionage (see Miller [30]).

[4]For instance, Jehiel et al. [25] show that non-trivial regret-free implementation is possible only for a very small, degenerate set of mechanism design problems with *multivariate* private information and interdependent valuations. On the other hand, there are many instances not only of non-trivial, but of *efficient* regret-free implementation in more specific settings with interdependent valuations and *univariate* private information, particularly in the context of auctions (see, for instance, Crémer and McLean [11], Maskin [28], Dasgupta and Maskin [12], Eső and Maskin [13], Jehiel and Moldovanu [20], Bergemann and Välimäki [5], Perry and Reny [35], Krishna [26], Ausubel [3]). While non of these settings impose any budget-balance requirements, they also lack the contingent-payment feature of our problem, rendering direct comparisons impossible.

free trading mechanism. That our paper nevertheless arrives at largely negative results is caused by a *further* distinguishing feature of our merger setting: While the classical trade problem posits a publicly known outside option for the buyer (i.e., his utility if no trade occurs), it is natural in many merger settings to assume that *both* parties' outside options are subject to private information.

In a more immediately related paper, Brusco et al. [7] also investigate the role of asymmetric information in merger negotiations. In contrast to our paper, they search for *efficient* (Bayesian) rather than regret-free mechanisms. Differences in the set-up aside,[5] our findings complement each other by revealing a strong overlap in the attainability of these two goals: First, both papers show that when private knowledge does not pertain to outside options, both efficiency and regret-freeness are attainable. Second, for a large class of remaining settings, our papers combine to show that *neither* goal is attainable.

Interestingly, there has recently been quite some investigation into the efficient *dissolution* of partnerships under interdependent valuations, which in some ways represents the inverse to the problem considered in this paper. This literature works out conditions under which it is possible to efficiently dissolve joint ownership of a firm when valuations for that firm differ and are privately known.[6] In contrast to our paper, whether it is efficient to dissolve or not is generally not an issue: The social value of the partnership is simply the share-weighted average of agents' individual valuations, so dissolution will be efficient whenever agents' valuations differ. Rather, the question is only one of allocating the single indivisible item owned by the partnership to the agent with the highest valuation. Thus, the problem is strongly related to auction design, with some additional difficulties caused by budget-balance requirements. In our setting, on the other hand, the main question is whether the partnership should be formed or not. In contrast, how shares in the

[5] On the one hand, the set-up in Brusco et al. [7] allows merger negotiations taking place with more than one potential partner. On the other hand, their setting is more restrictive regarding the structure imposed on profit functions' dependence on private information.

[6] See Fieseler et al. [14], Jehiel and Pauzner [23], Ornelas and Turner [34] and the survey in Moldovanu [31].

partnership are to be allocated in the event of a merger does not affect efficiency, as agents have pure common valuations concerning the partnership itself.

The paper is organized as follows. Section 2 introduces the basic model, describing both the merger environment and the mechanisms to be considered. Section 3 shows that if at least one party's outside option is common knowledge, efficient implementation is possible. Section 4 presents this paper's main result by showing that, for more generic environments, not only efficient but in fact any kind of non-trivial implementation is impossible. Providing a closer look at budget balance, Section 5 shows that the possibility of "burning money" off equilibrium reverses the largely negative results of Sections 4. Finally, Section 6 concludes.

2 The Model

This section presents the basis for our analysis of merger negotiations. Section 2.1 introduces the general setup and basic terminology. We motivate our framework with specific examples in Section 2.2.

2.1 Merger Environment and Mechanism

We consider merger mechanisms in an environment of the following type:

Definition 2.1. A *merger environment* $\mathcal{E}$ is a tuple $(T_1, T_2, \pi_1, \pi_2, \pi^{\mathrm{M}})$ with the following components:

(i) $T_i = [0, 1]$, $i = 1, 2$, is the *type space* for firm i;

(ii) $\pi_i : T_1 \times T_2 \to \mathbb{R}_+$ is the *stand-alone profit function* for firm i;

(iii) $\pi^{\mathrm{M}} : T_1 \times T_2 \to \mathbb{R}_+$ is the *merger profit function.*

Thus, an environment is essentially a description of how well each of the two parties would do on their own and what they could achieve together, taking account of the fact that each of these quantities can depend on each party's private information, where private information is represented by the joint type-space $\mathbf{T} \equiv T_1 \times T_2$. Agents each know their own type t_i ex-ante, but observe realized profits (π_i if no merger occurs, π^{M} if a merger occurs) only after the negotiation. For convenience, we will assume that π_1, π_2 and π^{M} are continuous.[7]

The dependence of (stand-alone) profit functions on information held by the *other* firm is a very natural feature of an oligopolistic environment, but none of our ensuing results depend on this (i.e., all results encompass the case in which $\frac{\partial}{\partial t_j}\pi_i \equiv 0$).

Finally, note that we do not model oligopolistic competition explicitly. Like most of the related literature, our analysis abstracts from actions taken by firms after the negotiation: Profits are fully determined by the type profile and whether or not the merger takes place.[8] Specifically, stand-alone profits are independent of any information revealed by firms during the negotiation process, so that signalling considerations neither play a role for firms' decision to participate in the mechanism nor for the behavior in the mechanism.

For the environment described in Definition 2.1, we investigate mechanisms which use a system of transfers and—provided a merger occurs—allocations of the merged entity's profits. Anticipating the usual revelation argument, we restrict our attention to *direct* merger mechanisms, for which each agent's report is restricted to his type space:

Definition 2.2. A *(direct) merger mechanism* $\mathcal{M}$ is a tuple $(m, \hat{\pi}^{\mathrm{M}}, p^0)$ consisting of a merger decision function m, merger-profit sharing rules $\hat{\pi}^{\mathrm{M}} = (\hat{\pi}_1^{\mathrm{M}}, \hat{\pi}_2^{\mathrm{M}})$, and transfer functions $p^0 = (p_1^0, p_2^0)$, which are defined as follows:

(i) The *merger decision function* $m : \mathbf{T} \to \{0, 1\}$ maps a combination of reports $\tilde{\mathbf{t}} = (\tilde{t}_1, \tilde{t}_2)$ by players 1 and 2 about their type into a merger decision, with

[7] Moreover, note that the above definition assumes that each agent's private information is univariate. Given the largely negative results concerning efficient implementation under interdependent valuations with *multi-dimensional* as opposed to *univariate* information in auction settings (Jehiel and Moldovanu [20]) and in more general settings (Jehiel et al. [25]), scalar-valued information should improve the scope for efficient implementation in our merger setting.

[8] For instance, Jehiel and Moldovanu [22, 19] and Jehiel et al. [24] consider auctions where payoffs reflect product market profits from an unmodeled game. In the Conclusion, we will discuss to which extent it is possible to extend our analysis to a setting where actions in the product market are modeled explicitly.

$m(\tilde{\mathbf{t}}) = 1$ if and only if the merger takes place as a result of the reports. We let $M^0 \equiv m^{-1}(0)$ and $M^1 \equiv m^{-1}(1)$.

(ii) The *transfer functions* $p_i^0 : M^0 \to \mathbb{R}$, $i \in \{1, 2\}$, specify a transfer payment from player i to the mechanism operator for any reported types $\tilde{\mathbf{t}} \in M^0$.

(iii) The *merger-profit sharing rules* $\hat{\pi}_i^{\mathrm{M}} : M^1 \times \mathbb{R}_+ \to \mathbb{R}$, $i \in \{1, 2\}$, map reports $\tilde{\mathbf{t}} \in M^1$ and realized merger profits $\pi^{\mathrm{M}}(\mathbf{t})$ into a payoff to firm i.

Thus, a merger mechanism's merger decision m produces a partition of type space $\mathbf{T}$ into type profiles M^0 such that no merger occurs (the 'no-merger set') and type profiles M^1 such that a merger occurs (the 'merger set'). For simplicity, we consider only merger mechanisms that are *well-behaved* in the sense that (i) the boundary between M^0 and M^1 is almost everywhere smooth, and (ii) for every $\mathbf{t} \in M^1$ and every ε-neighborhood $U_\varepsilon(\mathbf{t})$, there exists a point $\mathbf{t}' \in U_\varepsilon(\mathbf{t})$ in the interior of M^1.[9]

Letting $u_i(\tilde{\mathbf{t}}; \mathbf{t})$ denote agent i's payoffs for any combination of reported and true types $\tilde{\mathbf{t}}, \mathbf{t} \in \mathbf{T}$,

$$u_i(\tilde{\mathbf{t}}; \mathbf{t}) = \begin{cases} \pi_i(\mathbf{t}) - p_i^0(\tilde{\mathbf{t}}), & \text{if } \tilde{\mathbf{t}} \in M^0, \\ \hat{\pi}_i^{\mathrm{M}}[\tilde{\mathbf{t}}; \pi^{\mathrm{M}}(\mathbf{t})], & \text{if } \tilde{\mathbf{t}} \in M^1. \end{cases} \tag{1}$$

In the simplest case, agents' payoffs over M^1, $\hat{\pi}_i^{\mathrm{M}}(\cdot)$, may be thought of as arising through (report-dependent) allocations of shares in the merged entity and accompanying payments. More generally, the use of more sophisticated securities will permit an arbitrary conditioning of payoffs on reports and true post-merger profits.[10]

For reasons given in the introduction, we shall confine ourselves to finding *ex-post* implementable mechanisms, that is, mechanisms for which truthful reporting is a best response to any given type of the competitor, *given* truthful reporting by the other type. Formally, letting $U_i(\tilde{t}_i; t_i, t_j) = u_i(\tilde{t}_i, t_j; t_i, t_j)$ denote firm i's payoff given truthful reporting by the other firm, the requirement of *ex-post* incentive compatibility can be compactly formulated as follows:

[9] This will, for instance, exclude merger sets that contain isolated lines or isolated points.

[10] Observe that the transfer function p_i^0 is only defined over M^0 since over M^1, transfers (i.e., components of the payoff that are independent of the realized level of π^{M}) may be included w.l.o.g. in the merger-profit sharing rule $\hat{\pi}^{\mathrm{M}}$.

Definition 2.3. A merger mechanism $\mathcal{M}$ is *(ex-post) incentive compatible (IC)* if

$$U_i(t_i; t_i, t_j) \geqslant U_i(\tilde{t}_i; t_i, t_j), \qquad \text{for all } i \in \{1,2\},\ \tilde{t}_i, t_i \in T_i,\ t_j \in T_j. \tag{2}$$

We shall call a merger decision function m *implementable* if there exist $\hat{\pi}^{\mathrm{M}}$ and p^0 such that the mechanism $(m, \hat{\pi}^{\mathrm{M}}, p^0)$ is incentive compatible.[11]

Efficiency of a merger decision function is defined as follows:

Definition 2.4. A merger decision function m is *efficient* if, for all $\mathbf{t} \in \mathbf{T}$,

$$m(\mathbf{t}) = \begin{cases} 0, & \text{if } \pi^{\mathrm{M}}(\mathbf{t}) < \pi_1(\mathbf{t}) + \pi_2(\mathbf{t}), \\ 1, & \text{if } \pi^{\mathrm{M}}(\mathbf{t}) > \pi_1(\mathbf{t}) + \pi_2(\mathbf{t}). \end{cases} \tag{3}$$

We call a merger mechanism $\mathcal{M} = (m, \hat{\pi}^{\mathrm{M}}, p^0)$ efficient if and only if m is efficient.

Since it is natural to assume that no third party should benefit from or subsidize the mechanism, we introduce the following condition:

Definition 2.5. A merger mechanism is *budget balanced (BB)* if, for all $\tilde{\mathbf{t}}, \mathbf{t} \in \mathbf{T}$,

$$u_1(\tilde{\mathbf{t}}; \mathbf{t}) + u_2(\tilde{\mathbf{t}}; \mathbf{t}) = \begin{cases} \pi_1(\mathbf{t}) + \pi_2(\mathbf{t}), & \text{if } \tilde{\mathbf{t}} \in M^0, \\ \pi^{\mathrm{M}}(\mathbf{t}), & \text{if } \tilde{\mathbf{t}} \in M^1. \end{cases} \tag{4}$$

Importantly, this budget balance condition is required to hold not only for truthful reports.[12]

Finally, letting $V_i(t_i, t_j) \equiv U_i(t_i; t_i, t_j)$ denote agent i's *equilibrium payoff function* (or *value function*), mechanisms in which agents will voluntarily participate must satisfy the following requirement:

Definition 2.6. A merger mechanism is *individually rational (IR)* if

$$V_i(\mathbf{t}) \geqslant \pi_i(\mathbf{t}), \qquad \text{for all } i \in \{1,2\},\ \mathbf{t} \in \mathbf{T}. \tag{5}$$

[11]We have motivated our use of this ex-post concept of implementation by the desire to avoid ex-post regret, which makes most sense if agents learn each others' types ex-post. Observe to this end that, typically, agents' observation of ex-post profits (merger profits or stand-alone profits) permits perfect ex-post inference on the other's type. This inference will be perfect whenever the observed profit is strictly monotone in the other's type.

[12]Such a condition has typically been invoked also in the aforementioned literature on partnership dissolution (cf. Jehiel and Pauzner [23], Ornelas and Turner [34]).

To understand this definition, recall that in our set-up neither the decision to participate in the mechanism nor the merger decision reveals any information that influences the strategic interaction between the two parties after the merger.[13] Hence, stand-alone profits are the same no matter whether one party decides not to participate or whether the mechanism is played out and results in a no-merger decision. Moreover, consistent with our notion of *ex-post* incentive compatibility, (IR) requires *ex-post* rationality in that agents' equilibrium payoffs must weakly exceed their outside option for any possible type of the other player (rather than merely in expectation).[14]

Finally, observe that there is one particular mechanism which—albeit typically being highly inefficient—satisfies all remaining constraints, namely the mechanism which reproduces the status quo by prescribing (i) never merge, and (ii) always require zero payments. Such a mechanism satisfies (IC), (BB), and (IR).

2.2 Examples of Merger Environments

We illustrate our set-up with a few exemplary merger environments. The examples offer different instances of what the relevant private information t_1, t_2 consists in and, correspondingly, how it relates to stand-alone and post-merger profits.

Example 1 (Synergy Potentials)**.** Suppose each firm holds private information which concerns the synergy potentials in the event of a merger, but which is relevant to neither firm's stand-alone operations, so that outside options are commonly known. This information may relate to the ease with which each firm's production equipment, its workforce, or its organizational structure (including sales operations, etc.) can be integrated into a joint entity. More generally, it may relate to the profitability of some new business venture which is feasible only to a merged entity. If a higher value of t_i signals higher synergy potentials, efficiency requires that some

[13]Using the terminology of Jehiel and Moldovanu [21], there are no 'informational externalities' of merger negotiations on the ensuing interaction. Hviid and Prendergast [18] investigate such externalities in the context of mergers with one-sided asymmetric information.

[14]Assuming that both agents participate for any type profile involves no loss of generality since the mechanism can always reproduce the same outcome as if agents did not participate.

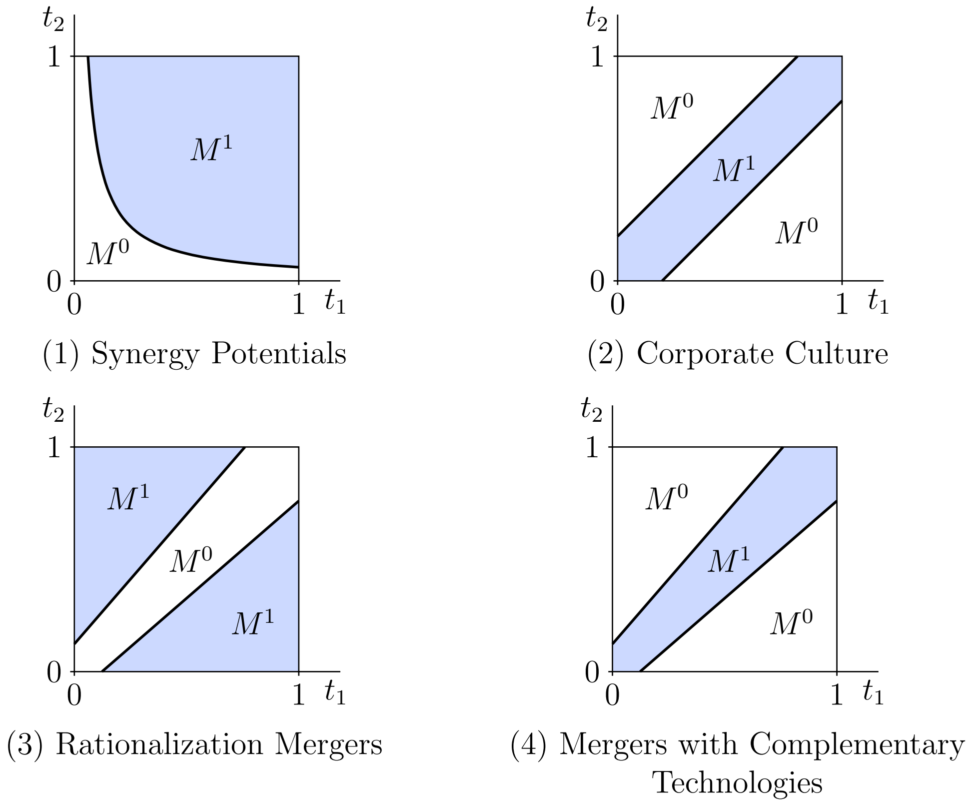

Figure 1: Efficient Merger Decisions for Examples of Section 2.2.

set of types above a downward-sloping line engages in a merger, as illustrated in panel (1) of Figure 1.

Example 2 (Corporate Culture)**.** As in the first example, suppose firms' private information is only relevant to post-merger profits. Suppose, however, that post-merger profits depend positively only on how *similar* firms' privately known types are. For instance, types could represent different patterns of corporate culture, all of which perform equally well in autarky, but where merging more dissimilar ones involves larger frictional losses.[15] Then the efficient set will typically contain points sufficiently close to the diagonal, as in panel (2) of Figure 1.

Examples 1 and 2 above have both assumed private information to be relevant

[15]Many case studies attribute failed mergers to incompatible corporate cultures of the merging firms; see Larsson and Finkelstein [27].

only in the event of a merger. However, if a firm's private information concerns factors such as production costs, workforce quality, financial health, demand forecasts, etc., this information will typically be relevant both to profits in the event of a merger *and* to the firm's stand-alone profits if no merger occurs. Moreover, provided some form of interaction between the firms—such as if firms operate on related markets—each firm's private information is relevant also to the competitor's profit if no merger occurs. Adding mild further conditions on the *direction* in which information affects each profit function, we shall call such a merger environment 'competitive:'

Definition 2.7. A merger environment is *competitive* if (i) $\partial\pi_i/\partial t_i > 0$, (ii) $\partial\pi_i/\partial t_j \leqslant 0$, and (iii) $\partial\pi^{\mathrm{M}}/\partial t_i > 0$ for all $i, j \in \{1, 2\}$, $i \neq j$.

Thus, in a competitive merger environment, any good news held by a firm concerning its stand-alone profits translates into good news concerning the prospects of a merged firm, but will be (weakly) detrimental to the competitor if no merger occurs. In spite of this common structure, competitive merger environments can still produce very different efficient merger decisions, as the following two examples illustrate:

Example 3 (Rationalization Mergers)**.** Consider a competitive merger environment in which a firm's private information concerns its technological know-how. If this know-how is easily transferrable to the technologically inferior firm's production process in the event of a merger, then merging will typically be more efficient for more unequal types, where this 'rationalization-effect' is strongest. Hence, the efficient set will typically consist of points away from the diagonal, as in panel (3) of Figure 1.

Example 4 (Mergers with Complementary Technologies)**.** Assume next that, in contrast to Example 3, the more efficient firm's advantage is *only imperfectly* transferrable in the event of a merger (such as if its advantage derives, for instance, from a better-trained workforce). At the extreme, if the *inferior* firm's technology is decisive to the merged entity's profitability, which may happen when the assets brought into the relationship are highly complementary, merging will typically be more efficient when the difference in firms' efficiency levels is small, so that the efficiency loss incurred by the more efficient firm is small. As illustrated in panel (4) of Figure 1, qualitatively, the efficient merger decision will correspond to the inverse of that under rationalization. Finally, scenarios in between these two polar cases will produce essentially arbitrary other partitions of the type-space.

3 Implementation with a Known Outside Option

In this section, we restrict attention to environments where the outside option of at least one player is commonly known. Clearly, this setting contains Examples 1 and 2 of the last section as special cases, where both outside options were commonly known. The main result shows that efficient implementation is possible in such a setting:

Proposition 3.1. *If the merger environment is such that neither agent holds any private information on some agent j's stand-alone profits, so π_j is constant in $\mathbf{t}$ for some $j \in \{1,2\}$, then any efficient merger decision can be implemented by a mechanism satisfying (BB) and (IR) by agent $i \neq j$ obtaining the full merger profits in the event of a merger and paying π_j to agent j, and zero transfers if no merger occurs.*

Under such a mechanism, agent j will always obtain his stand-alone profit no matter what signal he sends, so that truth-telling will always be weakly optimal. Concerning agent i, observe that transfers and shares are unaffected by any deviations which leave the merger decision unaffected, so that we only need to show incentive compatibility for deviations which affect the merger decision. Given any $\mathbf{t}$, agent i's payoff will be $\pi^{\mathrm{M}}(\mathbf{t}) - \pi_j$ from merging, and $\pi_i(\mathbf{t})$ from not merging. Hence, agent i's gains to merging are $\pi^{\mathrm{M}}(\mathbf{t}) - \pi_i(\mathbf{t}) - \pi_j$, which will be positive for any type constellation such that merging is efficient, and negative for any type constellation such that merging is inefficient. Therefore, given efficiency of the merger decision function m, a type such that merging is efficient (given the other type) will lose by deviating from truth-telling, and so will a type such that no merger occurs. It is easily seen that the described mechanism satisfies (BB) and (IR).

Intuitively, the mechanism employed in Proposition 3.1 makes the party whose outside option is private knowledge the residual claimant, thereby perfectly aligning this party's incentives with the goal of efficiency. The other party in turn will obtain a payoff equal to its known stand-alone profit *irrespective* of reports (and merger decision), and will therefore have no incentive to misreport.[16]

A trivial corollary of Proposition 3.1 is that implementation of the efficient merger decision is also possible if *both* agents' stand-alone profits are common knowledge, that is, if each agent's private information pertains only to profits under a merger.[17] Thus, both the 'Synergy Potentials'- and 'Corporate Culture'-examples discussed in Section 2.2 permit implementation of the efficient decision.

In light of this result, it may appear puzzling that many mergers appear to turn sour because of rows over corporate culture. There are two possible responses. First, in many cases, the problems of integrating the cultures of two companies may have less to do with asymmetric information on corporate cultures themselves than with *symmetric* uncertainty concerning how these cultures will blend. For instance, it appears unlikely that the owners of AOL and Time Warner were fully unaware of differences in management style, dress policy or e-mail systems. Rather it would seem that both parties underestimated the frictions involved in reconciling the differences.[18] Second, obviously, even when asymmetric information about corporate culture is involved, this does not preclude the coexistence of private information concerning other properties of the firms, rendering Proposition 3.1 irrelevant.

[16] This essentially corresponds to the mechanism proposed by Crémer [10] in a setting where a 'target' firm with known outside option is up for sale to a number of potential acquirers whose outside option is also known, but who possess private information concerning joint profits. Crémer shows that—in contrast to a pure cash auction—the target can come arbitrarily close to full surplus extraction by a contingent value auction such that the target receives essentially full residual claims to the merged entity and buys out the winning bidder by a cash transfer.

[17] This corollary parallels Proposition 2 in Brusco et al. [7], who fittingly use the term 'acquisition mechanisms' for merger mechanisms of this type.

[18] See Albarran and Gormly [1] for an account of the AOL/Time Warner merger.

Finally, a weak point of the mechanism in Proposition 3.1 is that truth-telling is only *weakly* optimal for the party whose stand-alone profits are known. This is clearly an undesirable property, but not uncommon in related problems (see e.g. Mezzetti [29]). However, if *both* parties' stand-alone profits are known (as in Examples 1 and 2 above)—so private information concerns only merger profits π^{M}—then efficiency can be achieved by means of a mechanism satisfying (IC), (BB) and (IR) for which truth-telling is *strictly* optimal. The simplest way to achieve this is to split merger profits in proportion to firms' (constant) stand-alone profits, so that $\hat{\pi}_i^{\mathrm{M}}(\tilde{\mathbf{t}};\mathbf{t}) = [\pi_i/(\pi_1+\pi_2)]\cdot\pi^{\mathrm{M}}(\mathbf{t})$. More generally, by complementary use of transfers over M^1, arbitrary divisions of merger profits are implementable along with an efficient merger decision. As long as each agent receives a strictly positive share, truth-telling will be strictly optimal for such mechanisms.

It is instructive to relate Proposition 3.1 to the well-known impossibility result of Myerson and Satterthwaite [32]. To this end, interpret our 'merger environment' as a 'trade environment' for the bilateral sale of some indivisible good, with $m = 1$ if trade occurs. Interpret player 1 as buyer and player 2 as seller, with types t_i corresponding to valuations for the good. Let $\pi_1(\mathbf{t}) \equiv 0$ denote the buyer's outside option, let $\pi_2(\mathbf{t}) = t_2$ denote the seller's outside option, and let $\pi^{\mathrm{M}}(\mathbf{t}) = t_1$ denote the buyer's utility in the event of trade. This reinterpretation of our model corresponds exactly to the Myerson-Satterthwaite setting—the only difference being that the mechanism now has the additional possibility of conditioning both agents' payoffs on the effective value of trade (i.e., on the buyer's utility if trade occurs). By Proposition 3.1, this possibility entirely reverses Myerson and Satterthwaite's negative finding that there exists no efficient Bayesian mechanism: Efficient trade is now implementable under the much stronger requirements of ex-post incentive compatibility, ex-post individual rationality, and budget balance.[19] This is achieved by making the seller

[19] In fact, Proposition 3.1 trivially applies to more general trade settings in which the seller's valuation $\pi_2(\mathbf{t})$ and the buyer's valution $\pi^{\mathrm{M}}(\mathbf{t})$ are *arbitrary* functions of $\mathbf{t}$, thus allowing for the possibility of common valuations for the good (i.e., the seller having information on the value of the good to the buyer and vice versa).

residual claimant on the buyer's utility in the event of trade—a mode of trade which admittedly strains possibilities in a classical trade setting, but poses a natural possibility in a merger context.

As we will show in the next section, however, this contingent-payment feature of merger negotiations loses its power once *both* parties' outside options are subject to private information. Cast into the Myerson-Satterthwaite setting, this corresponds to situations in which the buyer's utility *if no trade occurs* is subject to private information.

4 Implementation with Unknown Outside Options

Contrasting Proposition 3.1, we will now show that, within the class of competitive merger environments specified in Definition 2.7, not only is implementing the efficient merger decision generally impossible, but that any kind of second-best solution can only implement rather trivial merger decision functions.

The result uses several intuitive ingredients. These in turn derive separately from (a) incentive constraints which inhibit misreports that have *no* impact on the merger decision (relative to truth-telling), and from (b) incentive constraints which inhibit misreports that *alter* the merger decision.

Essentially, the constraints in (a) imply that any information revealed by the agents can at best be used to decide whether or not a merger takes place, but not to determine how aggregate profits are to be split among the two parties. This is easily seen for type profiles such that no merger takes place: If equilibrium transfers were dependent on types over M^0, some agent could necessarily increase his transfer by misreporting, without reducing his stand-alone profit. For type profiles such that a merger takes place, on the other hand, the result derives from two sources: First, *selective* punishment of a party which deviates from truth-telling is made impossible by the fact that payoffs can be conditioned on true types only through π^M. Second, *collective* punishments for such inconsistent outcomes are precluded by (BB), which requires that the full merger profits be distributed to the merging parties.

In turn, the constraints in (b), by which agents should neither want to induce nor prevent a merger by misreporting, imply that for critical type profiles near the boundary of the merger set, the stand-alone payoff (net of transfers) and the payoff in case of a merger should be sufficiently close together, leading to a natural *continuity* restriction. Moreover, the constraints imply an alignment between the merger decision and agents' private returns to merging.

For *competitive* merger environments, the combination of these restrictions precludes implementation of any but trivial merger decisions: In such environments, agents' stand-alone payoffs respond to private information in an inherently opposed way, which—given the other restriction on equilibrium payoffs—makes it impossible to achieve the required alignment between the merger decision and agents' private returns to merging.

In the following, we formally derive the described restrictions on the mechanism in Sections 4.1–4.3. In Section 4.4, we then show how they dramatically reduce the scope for implementation in competitive merger environments.

4.1 Restrictions on the Mechanism over M^0

Local incentive constraints on the no-merger set immediately imply that transfers p_i^0 must be independent of player i's report. The simple intuition for this is that over M^0 the returns to changing one's report are independent of the true type profile due to the additive separability of agents' payoff functions in true and reported types. Moreover, (BB) then immediately implies that each agent's transfer must also be independent of the other's report. Taking these two arguments together, transfer functions must be constant over connected subsets of M^0. Moreover, it is straightforward to see that requiring (IR) in addition immediately implies that transfers are zero. The following lemma collects these results:

Lemma 4.1. *For any mechanism satisfying (IC) and (BB),*

(a) the vector of transfers p^0 must be constant in reports over any open connected subset of M^0;

(b) $p^0(\tilde{\mathbf{t}}) = \mathbf{0}$ *for any* $\tilde{\mathbf{t}} \in M^0$ *if (IR) is required in addition.*

Cast in terms of restrictions on agents' valuation functions V_i, (IC) and (BB) thus imply by part (a) of Lemma 4.1 that for any agent i, over any open connected subset of M^0, V_i equals π_i up to a constant. By part (b), imposing (IR) in addition implies equality also in the constant term.

4.2 Restrictions on the Mechanism over M^1

Next, we show that restrictions imposed by (IC) and (BB) essentially imply constance of value functions on level sets of the merger profit function over M^1. As mentioned already, the result rests on two key restrictions concerning the mechanism's response to deviations from truth-telling. First, as payoffs can be conditioned on true types only through π^M, outcomes which constitute an apparent deviation from truth-telling cannot entail a *selective* punishment of the party responsible for the deviation. Second, *collective* punishments for such inconsistent outcomes are precluded by (BB), which requires that the full merger profits be distributed to the merging parties, even if reports are inconsistent with the observed profits. The following proposition formulates this intuitive idea somewhat more precisely.

Proposition 4.2. *Any mechanism satisfying (IC) and (BB) must be such that, over the interior of* M^1*, each agent i's value function* V_i *is constant on each connected component of the intersection of any level set of* π^M *and the interior of* M^1*.*

Graphically speaking, each agent's value function V_i must thus be constant as we move along any π^{M}-level curve in M^1. The remainder of Section 4.2 provides the derivation of this Proposition.

4.2.1 A Simple Case

Proposition 4.2 reflects a basic separation problem which comes out most clearly in the following special environment. Suppose that, unlike in the competitive environment, there are two profiles $\mathbf{t}' = (t_i', t_j')$ and $\mathbf{t}'' = (t_i'', t_j')$ which differ only in the type of agent i and which both induce a merger with the same level of merger profits $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'')$. Then a merger mechanism that only conditions on true types via merger profits cannot distinguish the outcome from truthfully reporting $\mathbf{t}''$ from the outcome where true types are $\mathbf{t}'$ but agent i falsely reports t_i''. Thus, preventing such a false report requires

$$V_i(\mathbf{t}') \geqslant \hat{\pi}_i^{\mathrm{M}}\big[t_i'', t_j'; \pi^{\mathrm{M}}(\mathbf{t}')\big] = V_i(\mathbf{t}''). \tag{6}$$

By the same token $V_i(\mathbf{t}'') \geqslant V_i(\mathbf{t}')$. Hence, $V_i(\mathbf{t}'') = V_i(\mathbf{t}')$, that is, agent i must obtain the same equilibrium payoff for both profiles $\mathbf{t}'$ and $\mathbf{t}''$.

4.2.2 Proving Proposition 4.2

The problem with the above restriction on value functions is not only that it applies only to special environments. In addition, it exclusively concerns bundles that differ only in one player's type. Invoking (BB) in addition to (IC), however, we can arrive at a result that applies also to type profiles $\mathbf{t}'$ and $\mathbf{t}''$ which differ in both types. Formally, this generalization consists in replacing the requirement that the profiles differ only in one agent's type with the much weaker requirement that both profiles are in the merger set. To develop this result, we first extend condition (6) above to this situation.

Lemma 4.3. *Under any mechanism satisfying (IC) and (BB), for any* $i \neq j \in \{1,2\}$ *and* $\mathbf{t}' \equiv (t_i', t_j')$, $\mathbf{t}'' \equiv (t_i'', t_j'') \in \mathbf{T}$ *such that* $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'')$ *and* $(t_i'', t_j') \in M^1$,

(a) $V_i(\mathbf{t}') + V_j(\mathbf{t}'') \geqslant \pi^{\mathrm{M}}(\mathbf{t}')$, *and*

(b) $V_i(\mathbf{t}') \geqslant V_i(\mathbf{t}'')$, *if* $\mathbf{t}'' \in M^1$, *in addition.*

The proof is provided in the Appendix.

To understand this result, consider the visualization provided in Figure 2, and let $\tilde{\pi}^{\mathrm{M}} \equiv \pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'')$. Now assume that the revelation game has resulted in reports (t_i'', t_j') and merger profits of $\tilde{\pi}^{\mathrm{M}}$. This outcome may either be the result of agent i having provided a false report given types $\mathbf{t}'$, or of agent j having provided a false report given types $\mathbf{t}''$.[20] Because the mechanism can condition payoffs only

[20] Generally, such an outcome may of course also result from both agents *simultaneously* deviating

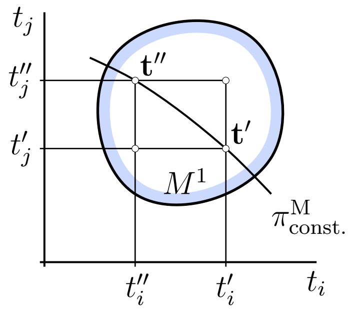

Figure 2: Illustration of Lemma 4.3 and Lemma 4.4.

on reports and observed merger profits, either situation must yield agents the same payoff.[21] Moreover, by (BB), agents' payoffs must sum to $\tilde{\pi}^{\mathrm{M}}$. Incentive compatibility (particularly, keeping i from reporting t''_i when true types are $\mathbf{t}'$ and agent j from reporting t'_j when true types are $\mathbf{t}''$) then immediately leads to Lemma 4.3(a). If $\mathbf{t}'' \in M^1$ in addition, then payoffs from truthful reporting in $\mathbf{t}''$ must sum to $\pi^{\mathrm{M}}(\mathbf{t}')$ as well. Thus, whatever agent j gets less from unilaterally misreporting t'_j when true types are $\mathbf{t}''$, agent i gets more, implying that agent i's payoff for reports (t''_i, t'_j) and true types $\mathbf{t}''$ must be at least $V_i(\mathbf{t}'')$, which produces part (b).

The next result is an immediate implication of Lemma 4.3 in situations where $\mathbf{t}'$, $\mathbf{t}''$, (t''_1, t'_2) and $t'_1, t''_2)$ all lie in M^1:

Lemma 4.4. *Under any mechanism satisfying (IC) and (BB), for any* $\mathbf{t}' = (t'_1, t'_2)$ *and* $\mathbf{t}'' = (t''_1, t''_2)$ *such that (i)* $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'')$ *and (ii)* $\mathbf{t}', \mathbf{t}'', (t''_1, t'_2), (t'_1, t''_2) \in M^1$, *we must have* $V_i(\mathbf{t}') = V_i(\mathbf{t}'')$, $i = 1, 2$.

This result extends the simple case of Section 4.2.1 to situations where *both* agents' types differ. Finally, Proposition 4.2, stated above, follows from iterated application of Lemma 4.4.[22]

from truth-telling. By the Nash-nature of *ex-post* incentive compatibility, however, only *unilateral* deviations from truth-telling must be made unprofitable.

[21]Intuitively, the two situations (i deviating or j deviating) are observationally equivalent to the mechanisms designer, i.e. he cannot verify *who* is responsible for the fraud or false report.

[22]For any two points $\mathbf{t}', \mathbf{t}''$ connected by a π^{M} level-curve lying in in the interior of M^1, a sequence of type profiles $\mathbf{t}^1, \ldots, \mathbf{t}^n$ can be found such that Lemma 4.4 is pairwise applicable to $(\mathbf{t}', \mathbf{t}^1), (\mathbf{t}^1, \mathbf{t}^2), \ldots, (\mathbf{t}^{n-1}, \mathbf{t}^n), (\mathbf{t}^n, \mathbf{t}'')$, implying $V_i(\mathbf{t}') = V_i(\mathbf{t}'')$, $i = 1, 2$.

4.3 Restrictions at the Boundary between M^0 and M^1

We now turn to deviations that affect the merger decision. We show that, to prevent such deviations, the mechanism must guarantee that (i) valuation functions are continuous at the boundary of the merger region and (ii) private incentives are aligned with the goals implicit in the merger decision.

To develop these results, we introduce the following terminology:

Definition 4.5. Agent i's type is *locally pivotal* to the merger decision m at $\mathbf{t} = (t_1, t_2) \in \mathbf{T}$ if and only if, for any $\varepsilon > 0$, there exists a t_i' with $|t_i - t_i'| < \varepsilon$ such that $m(t_i', t_j) \neq m(t_i, t_j)$.

At (almost) any point on the boundary, at least *one* agent's type is locally pivotal. Graphically, any parts of the boundary over which *only* agent i's type is locally pivotal will be perpendicular to the t_i-axis (i.e., either horizontal or vertical).

The two key restrictions to be developed in this section are (i) a continuity-property concerning agents' valuation functions across the merger boundary, and (ii) a condition ensuring that the change in the merger decision at the boundary is aligned with agents' fundamental preferences.

4.3.1 A Simple Case

As in Section 4.2, it is instructive to first develop a weaker version of these results in a simpler setting.

Proposition 4.6. *For any merger mechanism satisfying (IC), let $\hat{\mathbf{t}} = (\hat{t}_1, \hat{t}_2)$ denote any interior type profile such that some agent i's type is locally pivotal. Then, if π^{M} is constant in t_i over some neighborhood of $\hat{\mathbf{t}}$, the following must hold:*

(i) V_i is continuous in t_i at $\hat{\mathbf{t}}$;

(ii) π_i is weakly decreasing (increasing) in t_i at $\hat{\mathbf{t}}$ if m is increasing (decreasing).

To see (i), first note the following implication of the assumed local constance of π^M in t_i together with Lemma 4.1 and Proposition 4.2: There are constants $\tilde{V}_i$

and $\tilde{p}_i^0$, such that, for truthful reports t_j, for $j \neq i$, agents of type t_i earn $\tilde{V}_i$ for arbitrary reports near $\hat{t}_i$ that lead to a merger and $\pi_i(t_i, t_j) - \tilde{p}_i^0$ for arbitrary reports near $\hat{t}_i$ that do not. Types t_i such that truthful reporting leads to a merger will therefore misreport their type to prevent the merger unless $\tilde{V}_i \geqslant \pi_i(t_i, t_j) - \tilde{p}_i^0$. Similarly, types t_i such that truthful reporting leads to no merger will therefore misreport their type to induce a merger unless $\pi_i(t_i, t_j) - \tilde{p}_i^0 \geqslant \hat{\pi}_i^{\mathrm{M}}[(\tilde{t}_i, t_j); \pi^M(t_i, t_j)] = \hat{\pi}_i^{\mathrm{M}}[(\tilde{t}_i, t_j); \pi^M(\tilde{t}_i, t_j)] = \tilde{V}_i$, where we have used the presumption that π^M is locally constant in t_i; by continuity of π_i therefore $\pi_i(\hat{t}_i, t_j) - \tilde{p}_i^0 \geqslant \tilde{V}_i$, and hence $\pi_i(\hat{t}_i, t_j) - \tilde{p}_i^0 = \tilde{V}_i$.

Part (ii) is a simple implication of the more general concept of 'Positive Association of Differences' (PAD) property (see Roberts [37]).[23]

4.3.2 A More General Result

Again, Proposition 4.6 in itself is of limited value as it requires merger profits to be locally independent of some agent i's type. To understand the problems involved in its generalization, note that if, say, π^{M} were strictly increasing in t_i instead, then the payoff $\hat{\pi}_i^{\mathrm{M}}[\tilde{t}_i; \pi^{\mathrm{M}}(t_i)]$ obtained by a type $t_i \in (\hat{t}_i, 1]$ from reporting $\tilde{t}_i \in [0, \hat{t}_i)$ so as to prevent the merger no longer relates to any of the incentive constraints for types $t_i \in [0, \hat{t}_i)$ since those types' payoffs will all be conditioned on other levels of merger profits π^{M}.

However, as above, a generalization can be obtained by instead appealing to (BB) to link these incentive constraints to incentive constraints faced by the *other* agent for type profiles which involve the *same* merger profits π^{M}. This leads to a continuity requirement which, instead, holds on π^{M}-level curves:

---

[23]Specifically, by (PAD), if a change in agent i's signal from t_i to t_i' makes some alternative A relatively more profitable than some other alternative B (gross of transfers), then an incentive compatible mechanism must respect this by not choosing A at signal t_i and B at signal t_i'.

Proposition 4.7. *Suppose π^{M} is strictly monotone in both types. Then, under any mechanism satisfying (IC) and (BB), at any point $\hat{\mathbf{t}}$ on any smooth part of the boundary, the following must hold:*

(a) If agent i's type is locally pivotal at $\hat{\mathbf{t}}$, then

(ai) V_i must be continuous at $\hat{\mathbf{t}}$ along the level curve of π^{M} through $\hat{\mathbf{t}}$;

(aii) π_i must be weakly increasing in the direction of the π^{M}-level curve through $\hat{\mathbf{t}}$ in which m is decreasing;

(b) π_1 and π_2 must be constant on any connected subset of the boundary along which π^{M} is constant.

See the Appendix for the proof.

Part (a) extends the two insights from the simple case.[24,25] As described, it derives from combining agents' incentive constraints for different type profiles yielding the same merger payoffs, therefore producing a continuity requirement which holds on π^{M}-level curves. Part (b) can be understood as an extension of (aii) to a special case: If the merger boundary locally coincides with a level curve of π^{M}, then both agents' types will be locally pivotal, but the merger decision function m will generally be *constant* along this π^{M}-level curve. As part (b) shows, valuation functions must in this case also be constant along the relevant π^{M}-level curve (by the continuity-property (ai) and by transfers being constant over M^0, valuation functions being constant is equivalent to stand-alone profits being constant).

Proposition 4.7 has far-reaching consequences concerning the potential locations

[24]The statements in Proposition 4.7 are restricted to smooth parts of the boundary to make the proof less tedious. This will be sufficient for our purposes, however, given that we have assumed the boundary to be almost everywhere smooth.

[25]Note that by Proposition 4.6, Proposition 4.7(a) trivially extends also to merger environments where π^{M} is constant in *both* types. It runs into limitations, however, if π^{M} depends on *one* agent's type alone. To understand why, note that the underlying argument relies on the observed outcome of the unilateral deviation by agent i being replicable by some unilateral deviation on behalf of the other agent j from some other type profile, which is not possible if π^{M} depends on t_i alone. In that case, punishing agent i's deviation leads to no conflict with any of agent j's incentive constraints. Continuity need therefore no longer hold for this agent's value function.

of merger boundaries in type space. To see this, simply note that at points where continuity in *both* agents' value functions V_i is required, their sum must be continuous as well. Since $V_1 + V_2 \equiv \pi^{\mathrm{M}}$ over M^1 and $V_1 + V_2 \equiv \pi_1 + \pi_2$ over M^0 by (BB), this can only be the case at points $\mathbf{t}$ where $\pi^{\mathrm{M}}(\mathbf{t}) = \pi_1(\mathbf{t}) + \pi_2(\mathbf{t})$ (i.e., on the boundary of the efficient merger set). Moreover, the required alignment of the merger decision with individual preferences immediately translates into an alignment of the merger decision with *social* objectives:

Corollary 4.8. *Under the prerequisites of Proposition 4.7, any smooth parts of the boundary at which both agents' types are locally pivotal must be contained in the set* $\{\mathbf{t} \in \mathbf{T} \mid \pi^{\mathrm{M}}(\mathbf{t}) = \pi_1(\mathbf{t}) + \pi_2(\mathbf{t})\}$*. Moreover, at any such boundary point, m must be decreasing in any direction in which $\pi^{\mathrm{M}} - \pi_1 - \pi_2$ is strictly decreasing.*

See the Appendix for the proof.

Somewhat loosely speaking, except for horizontal or vertical parts of the boundary and parts where the boundary is not smooth, the boundaries of implementable merger decision functions must therefore be efficient in a local sense.

4.4 Implementation in a Competitive Merger Environment

Building on our previous results, we now argue that, within the class of *competitive* merger environments (see Definition 2.7), only merger decisions of a rather trivial nature can be implemented. Coarsely speaking, it is hard to avoid ex-post regret because parties fundamentally disagree about the circumstances in which they find staying alone attractive. More specifically, local changes in type profiles which leave joint merger profits constant will always *decrease* one agent's stand-alone profits. Given the constraints derived above, this precludes any merger decision which, for some type profile prescribing a merger, simultaneously gives *both* agents the possibility to inhibit the merger by a unilaterally false report.

To develop this result formally, we introduce the following definition:

Definition 4.9. A merger decision function m is *L-shaped* if, for some $\hat{\mathbf{t}} = (\hat{t}_1, \hat{t}_2) \in$

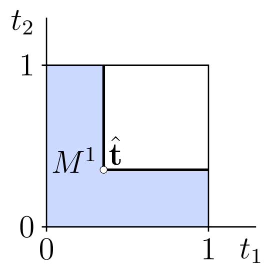

Figure 3: An L-Shaped Merger Decision Function.

$\mathbf{T}$ and any $\mathbf{t} = (t_1, t_2) \in \mathbf{T}$, it satisfies

$$m(\mathbf{t}) = \begin{cases} 0, & t_1 > \hat{t}_1 \text{ and } t_2 > \hat{t}_2, \\ 1, & t_1 < \hat{t}_1 \text{ or } t_2 < \hat{t}_2. \end{cases} \tag{7}$$

Figure 3 provides an illustration of an L-shaped merger decision function. Degenerate cases of an L-shaped merger decision function include mergers *never* occurring if $\hat{\mathbf{t}} = (0, 0)$, mergers *always* occurring if $\hat{\mathbf{t}} = (1, 1)$, and the merger decision globally depending on only one agent's report if $\hat{t}_1 = 0$ or $\hat{t}_2 = 0$.

The characteristic features of an L-shaped merger decision function are the following: (i) the merger decision is non-increasing in each type (as required for competitive merger environments by Proposition 4.7(aii)), and (ii) at each type profile in M^1, at most one agent can inhibit the merger by a false report. As the next result shows, any other merger decision function is not implementable under (BB):

Proposition 4.10. *In any competitive merger environment, under (BB), any merger decision function which is not L-shaped is not implementable.*

See the Appendix for the proof.

The intuition for Proposition 4.10 essentially combines two insights: First, it is easily seen that competitive merger environments have the property that players' stand-alone profits always decrease in *different* directions along any π^{M}-level curve. Second, however, aligning *both* players' private interests with the merger decision function requires *both* players' stand-alone profits to be locally decreasing along the

π^{M}-level curve wherever the merger decision is increasing (see Proposition 4.7(b)). This leaves only merger decision functions which are L-shaped as candidates for implementation—these decision functions being the only ones which require alignment of only *one* player's interests with the decision function at any boundary point.[26,27]

If we add (IR) to the requirements of Proposition 4.10, this immediately implies that $\pi^{\mathrm{M}}(\mathbf{t}) \geqslant \pi_1(\mathbf{t}) + \pi_2(\mathbf{t})$ for all $\mathbf{t} \in M^1$: If joint merger profits are lower than aggregate stand-alone payoffs, at least one agent must necessarily lose from the merger.[28] Combined with Proposition 4.10, the following result is immediate:

Corollary 4.11. *Unless a competitive merger environment is such that some agent i's type being low enough is a sufficient indicator that a merger is efficient, only trivial merger decisions of the type 'never merge' are implementable if the mechanism is to satisfy (IR) in addition to (BB).*[29]

[26]Conversely, however, L-shapedness is not a sufficient property to guarantee implementability of a merger decision function under (BB). However, it is easily shown that a weaker version holds: L-shaped merger decision functions for which $\hat{t}_i = 0$ for some $i \in \{1, 2\}$ are always implementable under (BB). These are degenerate instances of L-shaped merger decision functions, however, as at least one agent's report will be altogether irrelevant to the merger decision. More generally, it can be shown that an L-shaped merger decision function with vertex $\hat{\mathbf{t}} = (\hat{t}_1, \hat{t}_2)$, $\hat{t}_1, \hat{t}_2 > 0$, is implementable under (BB) if and only if for any $\mathbf{t}', \mathbf{t}'' \in M^1$ such that (i) $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'')$, and (ii) $t_1' = \hat{t}_1$ and $t_2'' = \hat{t}_2$, it holds that $\pi^{\mathrm{M}}(\mathbf{t}') \geqslant \max\{\pi_1(\mathbf{t}') + \pi_2(\mathbf{t}''), \pi_1(\mathbf{t}'') + \pi_2(\mathbf{t}')\}$.

[27]Proposition 4.10 is related to results derived by Hagerty and Rogerson [15]. In the context of bilateral trade à la Myerson and Satterthwaite [32], these authors show that the only ex-post implementable trading mechanisms are 'posted price' mechanisms. For such mechanisms in turn, the boundary of the area where trade occurs is similarly L-shaped. This similarity in the set of implementable allocations derives from the fact that, as in our model, private information can be used only to determine whether trade occurs or not, but not to determine how the surplus is shared in the event of trade. As pointed out in Section 3, however, the rationale for *why* private information can only be used in this limited way differs between the bilateral-trade setting and our merger setting.

[28]There is another argument in favor of *a priori* focussing attention to mechanisms which satisfy this latter condition, based on a renegotiation-proofness argument: Any time a mechanism produces an inefficient merger decision and parties have somehow revealed their type through the mechanism, there is ample scope for renegotiation. The applicability of any inefficient mechanism therefore relies crucially on the mechanism operator's ability to commit. However, standard reputation arguments to justify the credibility of such a commitment are likely to fail in the merger context, where the problem to be solved by the mechanism is of an inherent one-shot nature.

[29]We should point out, however, that adding the requirement $\pi^{\mathrm{M}}(\mathbf{t}) \geqslant \pi_1(\mathbf{t}) + \pi_2(\mathbf{t})$ for all $\mathbf{t} \in M^1$ to the conditions imposed on m in Proposition 4.10 is *not* sufficient to ensure that m is implementable under (IC), (IR) and (BB). This has to do with the fact that, whereas equilibrium

Thus, in a wide class of competitive merger environments (including Examples 3 and 4 above), a merger mechanism satisfying (BB) and (IR) can achieve no more than reproduce the status quo.

It is important to note that our definition of a competitive merger environment includes the degenerate case that firms' stand-alone profits are independent of the other firm's type (condition (ii) of Definition 2.7 holds with equality). Thus, the negative results of this section also hold in this seemingly simpler case.

4.5 Observable Stand-Alone Profits

Our analysis above has assumed that, while the profits π^{M} of a merged firm are ex-post observable and contractible,[30] individual stand-alone profits π_i if *no* merger occurs are not, implying that the mechanism can condition utility transfers between firms on ex-post outcomes only if a merger occurs.

There are clear reasons why firms might not be able to contract arbitrarily on stand-alone profits. Most notably, anti-trust legislation puts severe limitations on such arrangements. Nevertheless, to some extent, firms can clearly use arrangements such as cross-ownership (or stock transfers), which condition on stand-alone profits. With the following class of mechanisms, we therefore consider the opposite polar case of fully unconstrained use of ex-post information.

Definition 4.12. A *generalized (direct) merger mechanism with unconstrained use of ex-post information*, $\mathcal{M}^{\mathrm{U}}$, is a tuple $(m, \hat{\pi}^{\mathrm{M}}, \hat{\pi}^0)$ consisting of a merger decision function m, merger-profit sharing rules $\hat{\pi}^{\mathrm{M}} = (\hat{\pi}_1^{\mathrm{M}}, \hat{\pi}_2^{\mathrm{M}})$, and stand-alone profit shar-

utilities V_i must be constant along π^{M}-level curves over M^1, the outside options π_1 and π_2 will vary. More specifically, for any m described in Proposition 4.10 and any attainable level of merger profits $\tilde{\pi}^{\mathrm{M}}$, let $\mathbf{A}(\tilde{\pi}^{\mathrm{M}}) \equiv \{\mathbf{t} \in M^1 \mid \pi^{\mathrm{M}}(\mathbf{t}) = \tilde{\pi}^{\mathrm{M}}\}$ denote the intersection of the corresponding level curve with M^1, and let $\tilde{V}_1(\tilde{\pi}^{\mathrm{M}}), \tilde{V}_2(\tilde{\pi}^{\mathrm{M}})$ denote the (by Proposition 4.2 constant) equilibrium utility levels on $\mathbf{A}(\tilde{\pi}^{\mathrm{M}})$. Then (IR) requires $\tilde{V}_i(\tilde{\pi}^{\mathrm{M}}) \geqslant \max_{\mathbf{t}\in\mathbf{A}(\tilde{\pi}^{\mathrm{M}})} \pi_i(\mathbf{t})$, $i = 1, 2$, and hence $\tilde{\pi}^{\mathrm{M}} \geqslant \max_{\mathbf{t}\in\mathbf{A}(\tilde{\pi}^{\mathrm{M}})} \pi_1(\mathbf{t}) + \max_{\mathbf{t}\in\mathbf{A}(\tilde{\pi}^{\mathrm{M}})} \pi_2(\mathbf{t})$. It is quickly checked that this results in an upper bound on $\hat{t}_i$, the value of t_i at the vertex of the L-shaped merger decision function for any agent i (intuitively, so that the variation in π_1 and π_2 over M^1 is limited).

[30]Non-contractibility of stand-alone profits is perhaps the stronger concern: As long as π_i is strictly monotone in t_j, firm i can perfectly infer t_j (and thereby π_j) from observing its own stand-alone profits.

ing rules $\hat{\pi}^0 = (\hat{\pi}_1^0, \hat{\pi}_2^0)$ such that, for any report $\tilde{\mathbf{t}} \in \mathbf{T}$ and true types $\mathbf{t} \in \mathbf{T}$, agents' payoffs are

$$u_i(\tilde{\mathbf{t}}; \mathbf{t}) = \begin{cases} \hat{\pi}_i^{\mathrm{M}}[\hat{\mathbf{t}}; \pi^{\mathrm{M}}(\mathbf{t})], & \text{if } \tilde{\mathbf{t}} \in M^1, \text{ and} \\ \hat{\pi}_i^0[\tilde{\mathbf{t}}; \pi_i(\mathbf{t}), \pi_j(\mathbf{t})], & \text{if } \tilde{\mathbf{t}} \in M^0. \end{cases} \tag{8}$$

Thus, in contrast to our previous analysis, we now allow stand-alone payoffs after unsuccessful negotiations to depend on *both* firms' stand-alone payoffs in an arbitrary manner.

Since the (non-generalized) mechanisms considered previously are special cases of those given by Definition 4.12, the *positive* results obtained in Section 3 above are still valid within this wider class of mechanisms. As far as our negative results are concerned, recall that, in our above analysis, we had derived strong restrictions on implementable merger decision functions (particularly Proposition 4.10) using only (BB). Given this, (IR) played only a rather subordinate role.

This is different in the generalized class of mechanisms considered here. It is easily seen that, if agents can be *forced* to participate in the mechanism, then efficient implementation is possible with unconstrained use of ex-post information:

Lemma 4.13. *In any merger environment, any* efficient *merger decision can be implemented by a merger mechanism with unconstrained use of ex-post information which satisfies (BB).*

Proof. Simply set $\hat{\pi}_i^{\mathrm{M}}(\hat{\mathbf{t}}; \pi^{\mathrm{M}}) = \frac{1}{2} \cdot \pi^{\mathrm{M}}$ and $\hat{\pi}_i(\hat{\mathbf{t}}; \pi_i, \pi_j) = \frac{1}{2} \cdot (\pi_1 + \pi_2)$.[31] □

However, as we now argue, invoking (IR) generalizes the negative results of the above analysis to generalized merger mechanisms. The key to this lies in the following simple insight:

Lemma 4.14. *Any generalized mechanism satisfying (BB) and (IR) must be such that $V_i(\mathbf{t}) = \pi_i(\mathbf{t})$ for all $\mathbf{t} \in M^0$ and $i = 1, 2$.*

[31] Note, however, that this mechanism would most likely not be warmly received by competition authorities, as it effectively eliminates competition even in the absence of a merger.

The reason is simply that, due to (BB), agents' equilibrium payoffs over M^0 must sum to their outside options, so $V_1(\mathbf{t}) + V_2(\mathbf{t}) = \pi_1(\mathbf{t}) + \pi_2(\mathbf{t})$ for $\mathbf{t} \in M^0$. Since, by (IR), each agent's equilibrium payoff must satisfy $V_i(\mathbf{t}) \geqslant \pi_i(\mathbf{t})$, Lemma 4.14 immediately follows.

Consequently, as far as equilibrium outcomes are concerned, under (BB) and (IR), the unconstrained use of ex-post information does not improve the scope for implementation. Specifically, for each merger mechanism $\mathcal{M}^{\mathrm{U}} = (m, \hat{\pi}^{\mathrm{M}}, \hat{\pi}^0)$ with unconstrained use of ex-post information, we define the induced merger mechanism (with *constrained* use of ex-post information) $\mathcal{M}^{\mathrm{U|I}}$ as $(m, \hat{\pi}^{\mathrm{M}}, p_i^0)$ with $p_1^0(\tilde{\mathbf{t}}) = p_2^0(\tilde{\mathbf{t}}) = 0$, $\forall \tilde{\mathbf{t}} \in \mathbf{T}$. Thus $\mathcal{M}^{\mathrm{U|I}}$ has the same merger decision function and merger-profit sharing rules as $\mathcal{M}^{\mathrm{U}}$, but zero transfers over M^0.

The following proposition is the central result of this section:

Proposition 4.15. *Any merger decision function $m(\cdot)$ which is implementable with unconstrained use of ex-post information by $\mathcal{M}^{\mathrm{U}} = (m, \hat{\pi}^{\mathrm{M}}, \hat{\pi}_i^0)$ under (BB) and (IR) is implementable by the induced mechanism $\mathcal{M}^{\mathrm{U|I}}$.*

See the Appendix for the proof.

Proposition 4.15 implies that the negative results derived in Section 4.4 for standard merger mechanisms under (BB) apply also to mechanisms with unconstrained use of ex-post information if we require (IR) in addition. Particularly, Corollary 4.11 applies also to generalized merger mechanisms.

5 Relaxing Budget Balance

Our discussion in Section 4 has pointed out two important ingredients to the impossibility result developed there: *First*, that the mechanism typically cannot distinguish unilateral deviations by an agent over M^1 from deviations by the other, which precludes a selective punishment for deviations from truth-telling. *Second*, since (BB) requires budget balance also off-equilibrium, a collective punishment for outcomes which constitute an apparent deviation on some parties' behalf are also precluded.

This section picks up on the latter point by showing that the possibility of "burning money" off-equilibrium indeed makes efficient implementation possible in a large class of environments. To this end, we replace (BB), the budget balance concept used hitherto, with the following, weaker concept:

Definition 5.1. A merger mechanism satisfies *equilibrium budget balance (EBB)* if, for all $\mathbf{t} \in \mathbf{T}$,

$$V_1(\mathbf{t}) + V_2(\mathbf{t}) = \begin{cases} \pi_1(\mathbf{t}) + \pi_2(\mathbf{t}), & \text{if } \mathbf{t} \in M^0, \\ \pi^{\mathrm{M}}(\mathbf{t}), & \text{if } \mathbf{t} \in M^1. \end{cases} \tag{9}$$

In contrast to (BB), (EBB) allows the budget to be broken *off equilibrium*, that is, when agents misreport their types. This paves the way for efficient implementation in a large class of environments:

Proposition 5.2. *In any merger environment where π^{M} is strictly monotone in each type, any efficient merger decision can be implemented by a merger mechanism satisfying (IC), (IR), and (EBB).*

The formal proof, provided in the Appendix, confirms the obvious intuition that imposing a sufficiently harsh *collective* punishment for any outcomes involving inconsistencies in reports and realized profits will permit implementation of the efficient merger decision with essentially arbitrary profit-sharing rules.[32,33]

[32] Indeed, in keeping with this intuition, the mechanism employed in the proof of Proposition 5.2 employs a sufficiently large *uniform* punishment α_i for any merger outcome involving reports that are incompatible with observed profits (requirement (b) in the proof). However, a more differentiated punishment can reduce the extent of "off-equilibrium money-burning". To see this, note that for any possible off-equilibrium merger outcome consisting of reports $\tilde{\mathbf{t}} = (\tilde{t}_1, \tilde{t}_2)$ and observed profits $\pi^{\mathrm{M}} \neq \pi^{\mathrm{M}}(\tilde{\mathbf{t}})$, the assumed monotonicity of π^{M} implies that there exists at most one $\hat{t}_i \in T_i$ such that $\pi^{\mathrm{M}}(\hat{t}_i, \tilde{t}_j) = \pi^{\mathrm{M}}$, that is, such that the observed outcome is attainable by agent i unilaterally deviating from truthful reporting. If such a $\hat{t}_i$ exists, (IC) can thus be ensured by setting $\hat{\pi}_i^{\mathrm{M}}(\tilde{\mathbf{t}}, \pi^{\mathrm{M}}) = \pi_i(\hat{t}_i, \tilde{t}_j)$ for any such outcome (which constitutes the maximal off-equilibrium payoff which ensures (IC), i.e. the minimal extent of "money burning"); if no such $\hat{t}_i$ exists, $\hat{\pi}_i^{\mathrm{M}}(\tilde{\mathbf{t}}, \pi^{\mathrm{M}})$ may be set arbitrarily for this outcome.

[33] The monotonicity requirement in Proposition 5.2 is related to the separation problem discussed in Section 4.2.1, which arises even in the absence of off-equilibrium budget balance when π^{M} is not monotone. To understand this, suppose there exist $\mathbf{t}' = (t_i', t_j')$ and $\mathbf{t}'' = (t_i'', t_j'')$ with $t_j' = t_j''$, $t_i' \neq t_i''$, $\mathbf{t}', \mathbf{t}'' \in M^1$ and $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'')$. Then (IC) requires $V_i(\mathbf{t}') = V_i(\mathbf{t}'')$. This additional restriction will generally prohibit choosing $\hat{\pi}^{\mathrm{M}}$ such that $V_1(\mathbf{t}) \geqslant \pi_1(\mathbf{t})$ and $V_2(\mathbf{t}) \geqslant \pi_2(\mathbf{t})$ for $\mathbf{t} = \mathbf{t}', \mathbf{t}''$—at least so long as $\pi_1(\mathbf{t}') \neq \pi_1(\mathbf{t}'')$ and $\pi_2(\mathbf{t}') \neq \pi_2(\mathbf{t}'')$. Moreover, preventing false

Proposition 5.2's scope is of course limited by the extent to which "off-equilibrium money burning" poses a realistic possibility in merger negotiations. Depending on the merger environment, the amount of wealth destroyed can be severe, and absent an adequate commitment device, parties will have strong incentives to renege on such outcomes. On the other hand, given that possibilities are very limited *without* money burning by our results in Section 4, parties should have a strong interest to find a device which enforces such punishments.

6 Conclusion

This paper has shown that the potential for implementing merger decisions which avoid ex-post regret depends in an important way on the information concerning players' outside options. If at least one player's outside option is subject to neither party's private information, then efficient implementation is possible by means of a simple buyout of this player by the other party. When private information concerns both party's outside options, however, the scope for regret-free implementation of any non-trivial merger decisions is extremely limited, at least as long as budget balance is required also off equilibrium. Restricting the budget-balance requirement to hold only under truth-telling, however, restores efficient implementation quite generally, but requires a third-party intermediary appropriating parties' wealth in off-equilibrium outcomes. In sum, these findings indicate that in a wide range of situations, mergers under asymmetric information necessarily involve ex-post regret.

In the context of our non-existence result, one may of course wonder about the scope for implementing *randomized* merger decisions. While a full answer to this question is beyond the scope of this paper, two basic observations should be made: First, stochastic merger decisions will generally be inefficient. Regarding our main impossibility result, it trivially follows that *efficient* merger decisions are not implementable, even if we permit stochastic merger decisions. Second, stochastic

reports which lead to inefficient mergers requires that $V_i(\mathbf{t}'') \leqslant \pi_i(\mathbf{t}')$ for any $\mathbf{t}' \in M^0$ and $\mathbf{t}'' \in M^1$ (the rest as above), which collides with $V_i(\mathbf{t}'') \geqslant \pi_i(\mathbf{t}'')$ whenever $\pi_i(\mathbf{t}') \leqslant \pi_i(\mathbf{t}'')$.

merger decisions will generally involve ex-post regret in the sense that agents will regret their action given the final draw of nature. Consequently, stochastic merger decisions can only be *interim* regret-free.

A further limitation of our analysis lies in its reduced-form approach to stand-alone profits of non-merged firms, which ignores some potentially interesting signaling effects that arise if firms interact again after failed merger negotiations (for instance, as competitors on the product market): a firm's choice of actions at the negotiation stage will then influence its competitor's perception of its type, thereby its competitor's future action, and thereby in turn its own expected profits. This introduces a new strategic rationale not only into actions taken during the negotiation round itself, but also into firms' choice of whether to participate in negotiations in the first place, as the profits of a firm which refuses to participate in negotiations need no longer coincide with the profits if firms negotiate unsuccessfully.

The explicit consideration of such post-negotiation actions leaves our analysis unaffected if firms' optimal actions are independent of beliefs held about the other's type, as this eliminates the aforementioned signaling concerns. This in turn is a reasonable assumption in the case of conglomerate or cross-border mergers where unmerged firms are active in distinct markets, or in the 'Synergy Potentials'- and 'Corporate Culture'-examples of Section 2.2, where stand-alone profits are altogether independent of types. On the other hand, in cases where signaling does play a role, firms' actions are necessarily not ex-post optimal. Consequently, given that firms necessarily regret their post-merger actions, the no-regret requirement at the negotiation stage loses its appeal. Signaling effects would therefore most coherently be analyzed under an entirely Bayesian approach (i.e., requiring optimality of all actions only in *expected* terms), which we believe should represent an interesting avenue for future research.

In sum, our analysis has investigated the scope for regret-free merger negotiations in a setting which abstracts from various independent sources of ex-post regret, including exogenous sources of uncertainty, randomization in the negotiation process, and signaling effects. Given that ex-post regret is largely unavoidable even under these favorable circumstances, the natural next step should be to investigate the scope for implementation in the broader class of Bayesian mechanisms for which firms necessarily enter a gamble. Apart from the aforementioned extensions, such an investigation should be valuable in terms of uncovering the types of inefficiencies and regret caused by informational asymmetries and identifying possible 'lemons-market' properties of the merger market.

Appendix

Proof of Lemma 4.3. (a) To keep agent i from reporting t_i'' when true types are $\mathbf{t}'$ (and agent j reports truthfully), (IC) requires

$$V_i(\mathbf{t}') \geqslant \hat{\pi}_i^{\mathrm{M}}[t_i'', t_j'; \pi^{\mathrm{M}}(\mathbf{t}')]. \tag{A.1}$$

Similarly, to keep agent j from reporting t_j' when true types are $\mathbf{t}''$ (and agent i reports truthfully),

$$V_j(\mathbf{t}'') \geqslant \hat{\pi}_j^{\mathrm{M}}[t_j', t_i''; \pi^{\mathrm{M}}(\mathbf{t}'')]. \tag{A.2}$$

Since $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'')$ by presumption, the right-hand sides of (A.1) and (A.2) sum to $\pi^{\mathrm{M}}(\mathbf{t}')$ by (BB), which proves the claim.

(b) If $\mathbf{t}'' \in M^1$, then $V_j(\mathbf{t}'') + V_i(\mathbf{t}'') = \pi^{\mathrm{M}}(\mathbf{t}'')$ by (BB). Since the right-hand sides of (A.1) and (A.2) add to $\pi^{\mathrm{M}}(\mathbf{t}'')$ as well, (A.2) may be rewritten as $V_i(\mathbf{t}'') \leqslant \hat{\pi}_i^{\mathrm{M}}[t_i'', t_j'; \pi^{\mathrm{M}}(\mathbf{t}')]$, which together with (A.1) implies part (b). □

Proof of Proposition 4.7. Without loss of generality, we let π^{M} be strictly *increasing* in both types, so that the π^{M}-level curves are downward-sloping in t_1/t_2-space.[34] Further, due to symmetry, it suffices to prove the claim for agent 1.

Case A: Suppose first that the π^{M}-level curve and boundary intersect at $\hat{\mathbf{t}}$ and that the boundary has a strictly positive (possibly infinite) slope, as illustrated in Figure A.1(a). Then, any neighborhood of $\hat{\mathbf{t}}$ will contain $\mathbf{t}' \in M^0$, $\mathbf{t}'' \in M^1$ such that (i) $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'') = \pi^{\mathrm{M}}(\hat{\mathbf{t}})$, (ii) $(t_1'', t_2') \in M^1$.[35] Moreover, for any sufficiently small neighborhood of $\hat{\mathbf{t}}$, V_1 will be constant on the part of the π^{M}-level curve lying in M^1 by Proposition 4.2, and p_1^0 will be constant over M^0 by Lemma 4.1(a).[36] Restricting attention to such sufficiently small neighborhoods of $\hat{\mathbf{t}}$ and denoting these values by $\tilde{V}_1^1$ and $\tilde{p}_1^0$, respectively, Lemma 4.3(b) immediately implies $\pi_1(\mathbf{t}') - \tilde{p}_1^0 \geqslant \tilde{V}_1^1$. Continuity of π_1 thus implies

$$\pi_1(\hat{\mathbf{t}}) - \tilde{p}_1^0 \geqslant \tilde{V}_1^1. \tag{A.3}$$

[34] The other cases are obtained by a simple re-normalization of agents' type spaces.

[35] While Figure A.1(a) illustrates this for a merger decision function which is locally decreasing in agent 1's type, this is quickly seen to be true also in the locally increasing case.

[36] Larger neighborhoods might, for instance, contain parts of the level curve that belong to different connected components of the intersection with the interior of M^1.

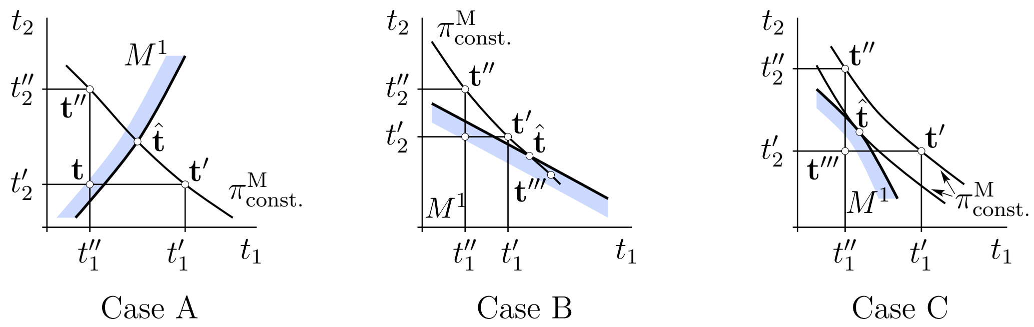

Figure A.1: Illustrations Accompanying Proof of Proposition 4.7.

Moreover, any such neighborhood of $\hat{\mathbf{t}}$ will contain a $\mathbf{t} \in M^1$ with $\pi^{\mathrm{M}}(\mathbf{t}) = \pi^{\mathrm{M}}(\hat{\mathbf{t}})$ such that agent 1 can inhibit the merger and obtain a payoff of $\pi_1(\mathbf{t}) - \tilde{p}_1^0$ (such as type profile $\mathbf{t}''$ in Figure A.1(a)), implying $\tilde{V}_1 \geqslant \pi_1(\mathbf{t}) - \tilde{p}_1^0$, and therefore

$$\tilde{V}_1 \geqslant \pi_1(\hat{\mathbf{t}}) - \tilde{p}_1^0. \tag{A.4}$$

(A.3) and (A.4) imply part (ai).

To see that part (aii) holds in this scenario, assume to the contrary that π_1 is *strictly decreasing* at $\hat{\mathbf{t}}$ in the direction along the part of the π^{M}-level curve running into M^0. Then, given the constance of V_1 over the part of the level curve lying in M^1 and the proven continuity result, there necessarily exist $\mathbf{t}' \in M^0$, $\mathbf{t}'' \in M^1$ with the same properties as above and such that $\pi_1(\mathbf{t}') - \tilde{p}_1^0 < \tilde{V}_1^1$, in contradiction to Lemma 4.3(b). Note that the above arguments carry through also if agent 2's report is not locally pivotal (i.e., if the boundary in Figure A.1(a) is vertical).

Case B: Next, consider a point of intersection $\hat{\mathbf{t}}$ where the boundary has a strictly negative slope, as in Figure A.1(b). Any neighborhood of $\hat{\mathbf{t}}$ will now contain $\mathbf{t}', \mathbf{t}'' \in M^0$ such that (i) $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'') = \pi^{\mathrm{M}}(\hat{\mathbf{t}})$ and (ii) either $(t_1'', t_2') \in M^1$ or $(t_1', t_2'') \in M^1$. Then Lemma 4.3(a) implies $\pi_1(\mathbf{t}') + \pi_2(\mathbf{t}'') \geqslant \pi^{\mathrm{M}}(\mathbf{t}')$. Again, since this is true of any sufficiently small neighborhood of $\hat{\mathbf{t}}$ and profit functions are continuous, it follows that $\pi_1(\hat{\mathbf{t}}) + \pi_2(\hat{\mathbf{t}}) \geqslant \pi^{\mathrm{M}}(\hat{\mathbf{t}})$. On the other hand, for any $\mathbf{t} \in M^1$ sufficiently close to $\hat{\mathbf{t}}$ and such that $\pi^{\mathrm{M}}(\mathbf{t}) = \pi^{\mathrm{M}}(\hat{\mathbf{t}})$ (such as $\mathbf{t}'''$ in Figure A.1(b)), either agent can inhibit the merger decision. Again restricting attention to small enough neighborhoods to let us use the notation $\tilde{V}_i^1$ and $\tilde{p}_i^0$ employed above, incentive compatibility requires that for any such $\mathbf{t}$

and any $i \in \{1, 2\}$,

$$\tilde{V}_i \geqslant \pi_i(\mathbf{t}) - \tilde{p}_i^0. \tag{A.5}$$

Therefore, summing the two incentive constraints, $\pi^{\mathrm{M}}(\hat{\mathbf{t}}) \geqslant \pi_1(\mathbf{t}) + \pi_2(\mathbf{t})$. Hence, $\pi^{\mathrm{M}}(\hat{\mathbf{t}}) \geqslant \pi_1(\hat{\mathbf{t}}) + \pi_2(\hat{\mathbf{t}})$, which, combined with our above result, immediately implies $\pi^{\mathrm{M}}(\hat{\mathbf{t}}) = \pi_1(\hat{\mathbf{t}}) + \pi_2(\hat{\mathbf{t}})$. Thus, $V_1 + V_2$ is continuous along the π^{M}-level curve.

By a limit argument, (A.5) must apply also at $\mathbf{t} = \hat{\mathbf{t}}$. However, given that $\pi^{\mathrm{M}}(\hat{\mathbf{t}}) = \pi_1(\hat{\mathbf{t}}) + \pi_2(\hat{\mathbf{t}})$, this condition can be met for both $i \in \{1, 2\}$ only if both conditions bind, which implies continuity of each V_i along the π^{M}-level curve. Proving part (aii) is straightforward by an argument as in Case A.

Case C: It thus remains to consider boundary points where boundary and π^{M}-level curve are tangent (this includes cases in which the boundary coincides with a π^{M}-level curve over some interval). Using a similar argument as in Case B, $\pi^{\mathrm{M}}(\hat{\mathbf{t}}) = \pi_1(\hat{\mathbf{t}}) + \pi_2(\hat{\mathbf{t}})$ at such a tangency point: Any neighborhood of $\hat{\mathbf{t}}$ contains $\mathbf{t}', \mathbf{t}'' \in M^0$ such that $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'')$ and such that $(t_1'', t_2') \in M^1$, implying $\pi^{\mathrm{M}}(\hat{\mathbf{t}}) \leqslant \pi_1(\hat{\mathbf{t}}) + \pi_2(\hat{\mathbf{t}})$ by Lemma 4.3(a) and a simple limit argument. Moreover, any neighborhood of $\hat{\mathbf{t}}$ contains a $\mathbf{t} \in M^1$ such that either agent can inhibit the merger with a unilateral misreport, implying by (BB) that $\pi^{\mathrm{M}}(\hat{\mathbf{t}}) \geqslant \pi_1(\hat{\mathbf{t}}) + \pi_2(\hat{\mathbf{t}})$.

Now let $(\mathbf{t}_n)$ denote any infinite sequence in M^1 which converges to $\hat{\mathbf{t}}$. Then, for n large enough (so that either agent can inhibit the merger), incentive compatibility combined with (BB) implies

$$\pi_1(\mathbf{t}_n) - \tilde{p}_1^0 \leqslant V_1(\mathbf{t}_n) \leqslant \pi^{\mathrm{M}}(\mathbf{t}_n) - \pi_2(\mathbf{t}_n) - \tilde{p}_1^0. \tag{A.6}$$

As $n \to \infty$, the left- and right-hand sides of (A.6) converge to $\pi_1(\hat{\mathbf{t}}) - \tilde{p}_1^0$ since $\pi^{\mathrm{M}}(\hat{\mathbf{t}}) = \pi_1(\hat{\mathbf{t}}) + \pi_2(\hat{\mathbf{t}})$ by our previous argument. Thus, V_1 is continuous at $\hat{\mathbf{t}}$, which completes the proof of part (a).

To prove part (b), simply note that for any two tangency points $\hat{\mathbf{t}}$ and $\hat{\mathbf{t}}'$ lying in a connected subset of the boundary on which π^{M} is constant, $\hat{V}_1^1(\hat{\mathbf{t}}) = \hat{V}_1^1(\hat{\mathbf{t}}')$: For any sequence $(\mathbf{t}_n)$ in M^1 converging to $\hat{\mathbf{t}}$, we can find a sequence $(\mathbf{t}_n')$ in M^1 which converges to $\hat{\mathbf{t}}'$ such that $\pi^{\mathrm{M}}(\mathbf{t}_n) = \pi^{\mathrm{M}}(\mathbf{t}_n')$ for every n and hence, by the constance of value functions over π^{M} level sets in M^1, $V_1(\mathbf{t}_n) = V_1(\mathbf{t}_n')$ for every n. Since $\hat{V}_1^1(\hat{\mathbf{t}}) = \pi_1(\hat{\mathbf{t}}) - \tilde{p}_1^0$ and $\hat{V}_1^1(\hat{\mathbf{t}}') = \pi_1(\hat{\mathbf{t}}') - \tilde{p}_1^0$, this immediately delivers $\pi_1(\hat{\mathbf{t}}) = \pi_1(\hat{\mathbf{t}}')$, as claimed. $\pi_2(\hat{\mathbf{t}}) = \pi_2(\hat{\mathbf{t}}')$

follows by symmetry. □

Proof of Corollary 4.8. The argument proving the first part of Corollary 4.8 is provided in the text. To see the second part, suppose to the contrary that there exists some direction in which $\pi^{\mathrm{M}} - \pi_1 - \pi_2$ is strictly decreasing but m is increasing at $\hat{\mathbf{t}}$. Then, any neighborhood of $\hat{\mathbf{t}}$ will contain a $\mathbf{t} \in M^1$ such that $\pi^{\mathrm{M}}(\mathbf{t}) < \pi_1(\mathbf{t}) + \pi_2(\mathbf{t})$ and such that both agents can inhibit the merger with some false report contained in this neighborhood. For sufficiently small neighborhoods, transfers are constant over M^0 in this neighborhood. Denoting these values by $\tilde{p}_i^0$, incentive compatibility requires $V_i(\mathbf{t}) \geqslant \pi_i(\mathbf{t}) - \tilde{p}_i^0$ and hence, by (BB), $V_1(\mathbf{t}) + V_2(\mathbf{t}) \geqslant \pi_1(\mathbf{t}) + \pi_2(\mathbf{t})$—a contradiction. □

Proof of Proposition 4.10. Step 1: Combining Proposition 4.7(ai) with the properties of competitive merger environments, it is immediately obvious that any smooth parts of the boundary must be either horizontal or vertical (see the argument in the main text). Thus, the merger boundary can consist only of combinations of horizontal and vertical segments.

Step 2: The next result restricts these cases further, leaving only the merger decision functions described in the proposition as potential candidates for implementation:

Lemma A.1. *In any competitive merger environment, any merger mechanism satisfying (IC) and (BB) must be such that,*

(a) at any smooth part of the merger boundary where only agent i's type is locally pivotal, the merger decision function m must be decreasing *in $\tilde{t}_i$;*

(b) the merger decision function exhibits no $\mathbf{t} \in M^1$ such that both *agents have the possibility of inhibiting the merger with some unilateral misreport.*

Part (a) is an immediate consequence of Proposition 4.7(aii): By the properties of competitive merger environments, π_i is strictly increasing on any π^{M}-level curve in the direction in which t_i is increasing, so that m must be decreasing in that direction.

To see part (b), observe first that, given step 1 of the argument, a merger set M^1 containing type profiles $\mathbf{t}$ such that *both* agents have the possibility of inhibiting the merger must (at least over some subset of type space) look as depicted in Figure A.2: an inverted L-shape with vertex $\hat{\mathbf{t}} = (\hat{t}_1, \hat{t}_2)$, where mergers occur if $\tilde{t}_1 < \hat{t}_1$ and $\tilde{t}_2 < \hat{t}_2$, and where no mergers occur if either $\tilde{t}_1 > \hat{t}_1$ or $\tilde{t}_2 > \hat{t}_2$. To see that such a merger

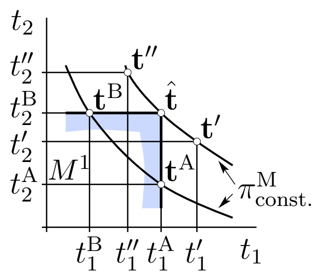

Figure A.2: Illustration Accompanying Proof of Proposition 4.10.

decision function cannot be implementable, observe first that, as shown in the figure, any neighborhood of $\hat{\mathbf{t}}$ will contain $\mathbf{t}', \mathbf{t}'' \in M^0$ lying on opposite sides of the π^{M}-level curve through $\hat{\mathbf{t}}$ such that $(t_1'', t_2') \in M^1$. Lemma 4.3(b) then implies $\pi^{\mathrm{M}}(\hat{\mathbf{t}}) \leqslant \pi_1(\mathbf{t}') + \pi_2(\mathbf{t}'')$, and hence, by a simple limit argument,

$$\pi^{\mathrm{M}}(\hat{\mathbf{t}}) \leqslant \pi_1(\hat{\mathbf{t}}) + \pi_2(\hat{\mathbf{t}}). \tag{A.7}$$

Moreover, as illustrated in Figure A.2, there exist $\mathbf{t}^{\mathrm{A}}, \mathbf{t}^{\mathrm{B}}$ on the boundary such that $t_1^{\mathrm{A}} = \hat{t}_1$, $t_2^{\mathrm{A}} < \hat{t}_2$, $t_2^{\mathrm{B}} = \hat{t}_2$, $t_1^{\mathrm{B}} < \hat{t}_1$, $\pi^{\mathrm{M}}(\mathbf{t}^{\mathrm{A}}) = \pi^{\mathrm{M}}(\mathbf{t}^{\mathrm{B}})$, and such that, by Proposition 4.2, value functions are constant on the π^{M}-level curve between $\mathbf{t}^{\mathrm{A}}$ and $\mathbf{t}^{\mathrm{B}}$. Denoting the latter values by $\tilde{V}_1$ and $\tilde{V}_2$, respectively, incentive compatibility for agent 1 at $\mathbf{t}^{\mathrm{A}}$ requires $\tilde{V}_1 \geqslant \pi_1(\mathbf{t}^{\mathrm{A}}) - \tilde{p}_1^0$, whereas at $\mathbf{t}^{\mathrm{B}}$, it requires $\tilde{V}_2 \geqslant \pi_2(\mathbf{t}^{\mathrm{B}}) - \tilde{p}_2^0$ for agent 2.[37] Thus,

$$\tilde{V}_1 + \tilde{V}_2 \geqslant \pi_1(\mathbf{t}^{\mathrm{A}}) + \pi_2(\mathbf{t}^{\mathrm{B}}). \tag{A.8}$$

But for competitive merger environments, $\pi_1(\mathbf{t}^{\mathrm{A}}) \geqslant \pi_1(\hat{\mathbf{t}})$ and $\pi_2(\mathbf{t}^{\mathrm{B}}) \geqslant \pi_2(\hat{\mathbf{t}})$ and $\tilde{V}_1 + \tilde{V}_2 < \pi^{\mathrm{M}}(\hat{\mathbf{t}})$ (due to π^{M} being increasing in both types and because of (BB)), implying

$$\pi^{\mathrm{M}}(\hat{\mathbf{t}}) > \pi_1(\hat{\mathbf{t}}) + \pi_2(\hat{\mathbf{t}}), \tag{A.9}$$

which contradicts (A.7) and thereby completes the proof of Lemma A.1.

Finally, given steps 1 and 2, only L-shaped merger decisions are left as possible candidates for implementation, which concludes the proof of Proposition 4.10. □

[37] In the by now familiar fashion, we restrict attention to sufficiently small neighborhoods of $\hat{\mathbf{t}}$ such that $p_i^0(\mathbf{t}) = \tilde{p}_i^0$ for all $\mathbf{t} \in M^0$ over this neighborhood. Moreover, note that strictly speaking, the above incentive constraints obtain directly only if the boundary points $\mathbf{t}^{\mathrm{A}}, \mathbf{t}^{\mathrm{B}}$ are contained in M^1. However, if this is not the case, the same restrictions obtain by a simple limit argument.

Proof of Proposition 4.15. Consider any merger mechanism with unconstrained use of ex-post information, $\mathcal{M}^{\mathrm{U}} = (m, \hat{\pi}^{\mathrm{M}}, \hat{\pi}^0)$, which satisfies (IC), (BB) and (IR). We will show that the induced mechanism $\mathcal{M}^{\mathrm{U|I}} = (m, \hat{\pi}^{\mathrm{M}}, p_i^0)$ with $p_1^0(\tilde{\mathbf{t}}) = p_2^0(\tilde{\mathbf{t}}) = 0$, $\forall \tilde{\mathbf{t}} \in \mathbf{T}$ then satisfies (IC) and (BB).

Concerning (BB), for reports $\tilde{\mathbf{t}} \in M^0$, the budget-balance condition is satisfied by construction of $p_i^0(\cdot)$. For reports $\tilde{\mathbf{t}} \in M^1$, on the other hand, the budget-balance conditions are identical for the two mechanisms (since merger-profit sharing rules $\tilde{\pi}_i^{\mathrm{M}}$ are identical). Thus, $\mathcal{M}^{\mathrm{U|I}}$ satisfies (BB).

Next, consider (IC). Recall that, in general terms, (IC) requires that, for any $i = 1, 2$, any $\mathbf{t} = (t_1, t_2) \in \mathbf{T}$, and any $\tilde{\mathbf{t}} = (\tilde{t}_1, \tilde{t}_2) \in \mathbf{T}$ such that $t_j = \tilde{t}_j$

$$u_i(\mathbf{t}, \mathbf{t}) \geqslant u_i(\tilde{\mathbf{t}}, \mathbf{t}), \tag{A.10}$$

is satisfied. To show that this holds for $\mathcal{M}^{\mathrm{U|I}}$, we will separately consider the four cases which result from the fact that $\mathbf{t}$ and $\tilde{\mathbf{t}}$ can each either lie in M^0 or in M^1.

For $\mathbf{t} \in M^1$ and $\tilde{\mathbf{t}} \in M^0$, $p_i^0 \equiv 0$ implies that (A.10) for $\mathcal{M}^{\mathrm{U|I}}$ is equivalent to $\hat{\pi}_i^{\mathrm{M}}[\mathbf{t}, \pi^{\mathrm{M}}(\mathbf{t})] \geqslant \pi_i(\mathbf{t})$, which is implied by the fact that $\mathcal{M}^{\mathrm{U}}$ satisfies (IR). For $\mathbf{t} \in M^0$ and $\tilde{\mathbf{t}} \in M^1$, (A.10) for $\mathcal{M}^{\mathrm{U|I}}$ is equivalent to

$$\pi_i(\mathbf{t}) \geqslant \hat{\pi}_i^{\mathrm{M}}[\tilde{\mathbf{t}}, \pi^{\mathrm{M}}(\mathbf{t})]. \tag{A.11}$$

Since, by Lemma 4.14, $V_i(\mathbf{t}) = u_i(\mathbf{t}, \mathbf{t}) = \pi_i(\mathbf{t})$ for $\mathbf{t} \in M^0$ for the mechanism $\mathcal{M}^{\mathrm{U}}$, (IC) and (BB) for $\mathcal{M}^{\mathrm{G}}$ implies (A.11). For $\mathbf{t}, \tilde{\mathbf{t}} \in M^0$, (A.10) is satisfied with equality since transfers p_i^0 are constant (at zero). Finally, for $\mathbf{t}, \tilde{\mathbf{t}} \in M^1$, incentive-compatibility conditions for $\mathcal{M}^{\mathrm{U}}$ and $\mathcal{M}^{\mathrm{U|I}}$ are equivalent (since merger-profit sharing rules $\hat{\pi}_i^{\mathrm{M}}$ are identical).

Hence, given that $\mathcal{M}^{\mathrm{U}}$ satisfies (IC), (BB) and (IR), $\mathcal{M}^{\mathrm{U|I}}$ satisfies (IC) and (BB), which completes the proof. □

Proof of Proposition 5.2. Choose the merger mechanism $(m, \hat{\pi}^{\mathrm{M}}, p_i^0)$ such that m is efficient, and such that

(a) $p_i^0(\tilde{\mathbf{t}}) = 0$ for $\tilde{\mathbf{t}} \in M^0$;

(b) $\hat{\pi}_i^{\mathrm{M}}(\tilde{\mathbf{t}}, \pi^{\mathrm{M}}) \equiv \alpha_i$, $i = 1, 2$, for all $\tilde{\mathbf{t}}, \pi^{\mathrm{M}}$ such that $\pi^{\mathrm{M}}(\tilde{\mathbf{t}}) \neq \pi^{\mathrm{M}}$, where α_i is a real number such that $\alpha_i < \inf_{\mathbf{t} \in \mathbf{T}} \pi_i(\mathbf{t})$;

(c) for $\tilde{\mathbf{t}}$ and π^{M} such that $\pi^{\mathrm{M}}(\tilde{\mathbf{t}}) = \pi^{\mathrm{M}}$, $\hat{\pi}_i^{\mathrm{M}}$ satisfies

(c1) $\hat{\pi}_i^{\mathrm{M}}(\tilde{\mathbf{t}}, \pi^{\mathrm{M}}) \geqslant \pi_i(\tilde{\mathbf{t}})$ for $i = 1, 2$, and

(c2) $\hat{\pi}_1^{\mathrm{M}}(\tilde{\mathbf{t}}, \pi^{\mathrm{M}}) + \hat{\pi}_2^{\mathrm{M}}(\tilde{\mathbf{t}}, \pi^{\mathrm{M}}) = \pi^{\mathrm{M}}(\tilde{\mathbf{t}})$ $(= \pi^{\mathrm{M}})$.

That is: a no-merger outcome involves no transfers by (a), type reports incompatible with observed profits under a merger entail sufficiently harsh punishments for both parties by (b), whereas equilibrium payoffs under a merger exceed standalone-profits for both parties by (c1) and add up to joint merger profits by (c2). Note that requirements (c1) and (c2) can jointly be met if the merger decision is efficient (more generally, if $M^1 \subseteq M^{1*}$).[38]

By requirements (a) and (c1), the mechanism satisfies (IR). (EBB) is trivially implied by (a) and (c2). As to (IC), consider first any $\mathbf{t}' = (t'_i, t'_j) \in M^0$ and consider i's incentives to report any $t''_i \neq t'_i$, leading to reports $\mathbf{t}'' \equiv (t''_i, t'_j)$. Since p_i^0 is constant in reports by requirement (a), such a deviation can only be profitable if $\mathbf{t}'' \in M^1$, that is, if it induces a merger. By requirement (b), however, this can only be profitable if $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'')$ (i.e., if the false report is compatible with realized merger profits), which in turn is precluded by joint merger profits π^{M} being strictly monotone in both types.

Next, consider truthful reports $\mathbf{t}' \in M^1$ and any unilateral deviation by agent i leading to reports $\mathbf{t}''$. Since $V_i(\mathbf{t}') = \hat{\pi}_i^{\mathrm{M}}[\mathbf{t}', \pi^{\mathrm{M}}(\mathbf{t}')] \geqslant \pi_i(\mathbf{t}')$ by (c1) and $p_i^0(\mathbf{t}'') = 0$ by (a), any such deviation involving $\mathbf{t}'' \in M^0$ cannot be profitable. For $\mathbf{t}'' \in M^1$ such that $\pi^{\mathrm{M}}(\mathbf{t}') \neq \pi^{\mathrm{M}}(\mathbf{t}'')$, requirements (b) and (c1) together imply $V_i(\mathbf{t}') \geqslant \hat{\pi}^{\mathrm{M}}[\mathbf{t}'', \pi^{\mathrm{M}}(\mathbf{t}')]$. A profitable deviation would thus need to involve $\pi^{\mathrm{M}}(\mathbf{t}') = \pi^{\mathrm{M}}(\mathbf{t}'')$, which is again precluded by π^{M} being strictly monotone in both types. □

[38] Also, note that the surplus generated by the merger may be split among parties in an arbitrary fashion.

References

[1] A. B. Albarran and E. K. Gormly, Strategic response or strategic blunder? An examination of AOL Time Warner and Vivendi Universal, in: R. G. Picard (Ed.), Strategic Responses to Media Market Changes, 2004, Jönköping International Business School Research Reports 2004–2.

[2] G. Andrade, M. Mitchell, and E. Stafford, New evidence and perspectives on mergers, J. Econ. Perspect. 15 (2001) 103–120.

[3] L. R. Ausubel, An efficient ascending-bid auction for multiple objects, Amer. Econ. Rev. 94 (2004) 1452–1475.

[4] D. Bergemann and S. Morris, Robust mechanism design, Econometrica 73 (2005) 1771–1813.

[5] D. Bergemann and J. Välimäki, Information acquisition and efficient mechanism design, Econometrica 70 (2002) 1007–1033.

[6] B. T. Brown, G. B. Murr, and D. A. Pluchinsky. Risky business. *Energy*, Winter 2005.

[7] S. Brusco, G. Lopomo, and S. Viswanathan, Efficient mechanisms for mergers and acquisitions, Int. Econ. Rev. 48 (2007) 995–1035.

[8] Business Week. Henry Silverman's long road back. February 28th, 2000, pp126–136.

[9] K. Chung and J. C. Ely. Ex-post incentive compatible mechanism design. Mimeo, Northwestern University, 2003.

[10] J. Crémer, Auctions with contingent payments: Comment, Amer. Econ. Rev. 77 (1987) 746.

[11] J. Crémer and R. P. McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent, Econometrica 53 (1985) 345–361.

[12] P. Dasgupta and E. Maskin, Efficient auctions, Quart. J. Econ. 115 (2000) 341–388.

[13] P. Eső and E. Maskin. Multi-good efficient auctions with multi-dimensional information. mimeo, Harvard University, 2000.

[14] K. Fieseler, T. Kittsteiner, and B. Moldovanu, Partnerships, lemons, and efficient trade, J. Econ. Theory 113 (2002) 223–234.

[15] K. M. Hagerty and W. P. Rogerson, Robust trading mechanisms, J. Econ. Theory 42 (1987) 94–107.

[16] R. G. Hansen, Auctions with contingent payments, Amer. Econ. Rev. 75 (1985) 862–865.

[17] P. Hietala, S. N. Kaplan, and D. T. Robinson, What is the price of hubris? Using takeover battles to infer overpayments and synergies, Finan. Manage. 32 (2003) 5–31.

[18] M. Hviid and C. Prendergast, Merger failure and merger profitability, J. Ind. Econ. 41 (1993) 371–386.

[19] P. Jehiel and B. Moldovanu, Auctions with downstream interaction among buyers, RAND J. Econ. 31 (2000) 768–791.

[20] P. Jehiel and B. Moldovanu, Efficient design with interdependent valuations, Econometrica 69 (2001) 1237–1259.

[21] P. Jehiel and B. Moldovanu, Allocative and informational externalities in auctions and related mechanisms, in: R. Blundell, W. Newey, and T. Persson (Eds.), The Proceedings of the 9th World Congress of the Econometric Society, Cambridge University Press, Cambridge, 2006.

[22] P. Jehiel and B. Moldovanu, Strategic nonparticipation, RAND J. Econ. 27 (1996) 84–98.

[23] P. Jehiel and A. Pauzner, Partnership dissolution with interdependent valuations, RAND J. Econ. 37 (2006) 1–22.

[24] P. Jehiel, B. Moldovanu, and E. Stacchetti, Multidimensional mechanism design for auctions with externalities, J. Econ. Theory 85 (1999) 258–293.

[25] P. Jehiel, M. Meyer-ter-Vehn, B. Moldovanu, and W. R. Zame, The limits of ex post implementation, Econometrica 74 (2006) 585–610.

[26] V. Krishna, Asymmetric english auctions, J. Econ. Theory 112 (2003) 261–288.

[27] R. Larsson and S. Finkelstein, Integrating strategic, organizational, and human resource perspectives on mergers and acquisitions: A case study of synergy realization, Organizational Science 10 (1999) 1–26.

[28] E. Maskin, Auctions and privatization, in: H. Siebert (Ed.), Privatization, Institut für Weltwirtschaft an der Universität Kiel, 1992, pp115-136.

[29] C. Mezzetti, Mechanism design with interdependent valuations: Efficiency, Econometrica 72 (2004) 1617–1626.

[30] D. A. Miller. The dynamic cost of ex post incentive compatibility in repeated games of private information. Mimeo, UCSD, 2007.

[31] B. Moldovanu, How to dissolve a partnership, J. Inst. Theoretical Econ. 158 (2002) 66–80.

[32] R. B. Myerson and M. A. Satterthwaite, Efficient mechanisms for bilateral trading, J. Econ. Theory 29 (1983) 265–281.

[33] M. Officer, Collars in merger negotations, J. Finance 59 (2004) 2719–2743.

[34] E. Ornelas and J. L. Turner. Efficient dissolutions of parterships and the structure of control. Mimeo, University of Georgia, 2004.

[35] M. Perry and P. Reny, An efficient auction, Econometrica 70 (2002) 1199–1212.

[36] J. G. Riley, Ex post information in auctions, Rev. Econ. Stud. 55 (1988) 409–430.

[37] K. W. Roberts, The characterization of implementable choice rules, in: J.-J. Laffont (Ed.), Aggregation and Revelation of Preferences, North-Holland, Amsterdam, 1979, pp321–348.

[38] The Economist. A Bavarian botch-up. August 3rd, 2000.