Profitable Double Marginalization*

Kevin A. Bryan†

Erik Hovenkamp‡

May 13, 2016

Abstract

When successive monopolies transact through noncooperative linear pricing, the resulting double markup decreases their joint profits relative to vertical integration. However, if there are downstream rivals (which are *not* double marginalized), the same noncooperative interaction often inadvertently raises their joint profits. Profit effects depend on how the well-understood harm from misaligned interests compares to the value of the resulting strategic effect. When profitable, vertical noncooperation incidentally approximates strategic delegation à la Bonanno and Vickers [1988], but avoids its credibility problem, suggesting an inability to bargain may be indirectly beneficial. The "conjectural consistency" concept helps to explain the disparate profit effects, and to synthesize the literature on strategic delegation and vertical control. The optimal way to "distort" a downstream firm's behavior is always to make it behave *as if* it has a consistent conjecture, no matter the distortion mechanism. If upstream competitors do this in parallel, they induce a "consistent conjectures equilibrium" (CCE) – or else a close analogue – evincing a strong link between ordinary Nash games and the CCE.

*We are grateful to Aviv Nevo, Robert Porter, Steve Salop, Glen Weyl, Heski Bar-Isaac, Kenneth Corts, and seminar participants at the University of Toronto and the University of Waterloo for their helpful feedback.

†University of Toronto, Rotman School of Management.

‡Northwestern University, Department of Economics and Pritzker School of Law. Erik's work on this project was supported in part by the Ewing Marion Kauffman Foundation.

1 Introduction

The double marginalization problem arises when producers of complementary goods set above-cost prices noncooperatively rather than jointly, leading the total price to reflect two or more independent markups. This problem was first recognized by Spengler [1950], and has played a prominent role in shaping antitrust and the theory of the firm.[1] In the case of successive monopolies, it is well-established that double marginalization both injures consumers and erodes joint profits relative to integration. Indeed, vertical merger is widely thought to enhance profits, since "integration harmonizes interests" (Williamson [1971]).

However, in oligopoly environments, pure integration may not maximize joint profits. Vickers [1985] shows that a firm may benefit by "strategically delegating" decisions to an agent with different objectives, for such delegation may act as a form of commitment.[2] For example, under Bertrand competition with differentiated products, a firm would like to commit to a higher price than it sets in the Bertrand-Nash equilibrium, and thus the firm could raise profits by delegating to an agent that always sets higher prices.[3] Beginning with Bonanno and Vickers [1988], a number of papers have applied the strategic delegation idea to vertical contracting, showing that double marginalization (or vertical restraints) could in principle be used *strategically* to increase joint profits.[4] This gives rise to the possibility of "strategic double marginalization," which involves a two-part tariff (whose marginal price could be higher or lower than marginal cost) set to maximize joint profits by exactly inducing the downstream firm to play its Stackelberg strategy in equilibrium – an idea that is by now well-understood.

Of course, this kind of strategic vertical contracting is very different from the noncooperative pricing that characterizes the successive monopolies problem; it is *designed* to maximize joint profits, whereas a noncooperative upstream firm disregards its impact on downstream profits. Furthermore, strategic delegation has been criticized for its apparent lack of credibility (e.g. Katz [1991]).[5] O'Brien and Shaffer [1992] specifically argue that two-part tariffs are non-credible to the extent that they can be privately renegotiated. Intuitively, if third party rivals are acting like Stackelberg followers, a vertical pair maximizing joint profits would like to deviate from the Stackelberg leader action (a Stackelberg leader's strategy is generally not a best response, after all.) Therefore, if non-public side agreements between the upstream and downstream firm are possible – or more generally

[1] See, e.g., Riordan [2008]; Salop and Scheffman [1983]; Williamson [1971].

[2] As was first established by Heinrich von Stackelberg, an oligopolist can often benefit by committing to a different strategy than it plays in equilibrium, because this allows it to take advantage of its rival's strategic behavior.

[3] See also, e.g., Fershtman and Judd [1987].

[4] See also Rey and Stiglitz [1995].

[5] For example, in the "divisionalization" model of Baye et al. [1986], distinct agents operate a number of competing intra-firm divisions (e.g. competing subsidiaries owned by a common parent) as a way to "precommit" to high quantity and hence cause rivals to act as Stackelberg followers. Once rivals have responded as such, the possibility of renegotiation causes divisionalization to unravel and become non-credible. After all, the parent company has both the power and the economic incentive to change the behavior of division heads once downstream rivals restrict production. Corts and Neher [2003] show that when a single upstream firm deals with multiple downstream firms, incentives for bilateral renegotiation, and hence the credibility of divisionalization, can be maintained. But if there but a single decision maker for a given firm (as in most of the delegation literature), the possibility of secret side agreements means that strategic delegation cannot be used profitably.

if the ability to contract efficiently implies the ability to renegotiate – then strategic delegation will tend to be non-credible and hence will not affect equilibrium outcomes.

On the other hand, what if a positive upstream markup results from an inability to coordinate price efficiently (as in the successive monopolies problem) rather than strategic contracting? The positive upstream markup will still affect the downstream firm's behavior, and in some environments this inadvertently improves joint profits by generating a valuable strategic effect, essentially approximating strategic delegation. But now there is no credibility problem, since the firms cannot coordinate in the first place. The result is that noncooperative double marginalization may provide a less precise but more credible way to generate the rents of strategic delegation.

In this paper, we provide a complete taxonomy of the situations in which double marginalization increases joint profits. We first explore *noncooperative* double marginalization in the spirit of Spengler [1950], where the upstream firm does not take into account the profits of the downstream firm, but rather charges a linear price to maximize upstream profits. Since we assume double marginalization does not directly affect any third party competitors, this treatment closely analogizes the successive monopolies problem. We then consider *cooperative* or *strategic* vertical contracting in the spirit of Bonanno and Vickers [1988], and we show how the notion of "conjectural consistency" serves to synthesize the literature on strategic delegation and vertical control. Our analysis is primarily directed at three questions.

1. Under what conditions does noncooperative linear pricing in a vertical relationship increase joint profits?

We study a single vertical relationship in which the upstream firm is a monopolist and the downstream firm faces third party competition (a single downstream rival).[6] Downstream competition may be in prices, quantities, or any other choice variable that would shape a firm's consumption of a costly input, e.g. R&D relying on a patented process. To that end, part of our contribution is to show that the profit effects of noncooperative double marginalization vary substantially among different downstream interactions. The downstream rival's costs are exogenously given and thus are not affected by the double marginalization problem; for example, it may be the the rival produces the input internally or simply does not use it.[7] In this sense, our model is essentially the simplest oligopoly extension of the successive monopolies problem: we add a downstream rival with exogenous input costs, but continue to focus on the same noncooperative interaction in the original vertical relationship.

That the downstream rival's costs are left exogenous distinguishes our analysis from some prior studies showing that noncooperative double marginalization can enhance joint profits when it is imposed on most or all downstream firms in parallel, generating a collusion-like effect in the downstream market (e.g. Rey and Stiglitz [1995]; Gaudet and Long [1996]) – or, relatedly, when it

[6] The results extend in a straightforward way to the case of multiple downstream rivals by transforming the interaction into an aggregate game.

[7] The point here is to isolate double marginalization to a single vertical chain when performing the relevant comparative statics – the comparison of joint profits between vertical integration and double marginalization.

can be used as a tool for raising rivals' costs (Gaudet and Long [1996], Salop and Scheffman [1983]). Unlike these studies, our model isolates all cost distortions created by double marginalization to the vertical pair whose joint profits are the focus of the comparative statics, and hence our results are not driven by any collusion-like or exclusionary effects.

We show that double marginalization inadvertently increases joint profits in many familiar environments, even though the upstream firm charges a monopoly price, and even though it does not directly affect the downstream rival. The noncooperative upstream markup depends only on equilibrium strategies, which shape demand elasticity for the upstream firm's good, whereas the jointly-optimal upstream markup depends on the effect of a downstream firm's action on the induced best response of rivals ("strategic effects"). These two things are largely independent of one another, however, because best response functions (which determine equilibrium strategies) ignore strategic effects by definition. Consequently, the impact of double marginalization on joint profits is entirely ambiguous: it can diminish, raise, or even maximize joint profits, compared to those of an integrated firm; in some games with strong strategic effects, the upstream monopoly price is too *low* to maximize joint profits. To the extent that *strategic* double marginalization is non-credible, our results suggest that firms with complementary interests may actually benefit from the inability to contract efficiently.

Section 2 provides examples of profitable double marginalization within familiar competitive environments, and also shows that our results carry over to the case of "horizontal" double marginalization. No single superficial property of the downstream game - such as whether it involves strategic substitutes or strategic complements - is sufficient to determine how double marginalization affects profits, because no such property entirely determines the impact of strategic effects on joint profits. In Section 3 we derive our primary results on the profit effects of double marginalization using a very general game that nests virtually all differentiable and concave models of simultaneous-move competition.

2. What explains the disparate profit effects of noncooperative double marginalization within different competitive environments?

Section 4 seeks to provide a richer understanding of why double marginalization produces such disparate profit effects within different games. We answer this question by invoking the theory of conjectural variation, first introduced by Bowley [1924], and in particular the notion of conjectural consistency popularized by Bresnahan [1981]. A firm's "conjecture" is its implicit belief about how rivals will respond to changes in its own conduct. In most games, firms' conjectures are "inconsistent," i.e. incorrect.[8] For example, a Cournot firm fashions a best response under the implicit belief (its conjecture) that rival output is fixed, but in truth the rival has a downward sloping reaction function. A Cournot firm with a consistent conjecture would internalize strategic effects, leading it to play the Stackelberg leader strategy.

[8]Specifically, a conjecture is a belief about the slope of one's rival's reaction function, and it is consistent if and only if this belief is correct, at least locally.

Depending on how a firm's inconsistent conjecture compares to the consistent one, the firm either understates or overstates the competitiveness of its rival's behavior. Noncooperative double marginalization can enhance joint profits in precisely those games in which the downstream firm's inconsistent conjecture leads it to overstate competition, in which case it can benefit from being induced to behave less competitively by paying a positive transfer to the upstream firm. Further, in every game, the jointly-optimal upstream markup (which could be negative) is always that which induces the downstream firm to behave *as if* it has a consistent conjecture, at least in equilibrium. One intuitive explanation is that an inconsistent conjecture reflects a failure to internalize strategic effects, and the jointly-optimal upstream markup acts like a Pigouvian tax that forces the downstream firm to internalize them.

3. How does strategic vertical pricing relate to other methods of distorting behavior in vertical relationships?

Beginning with Fershtman and Judd [1987] and Bonanno and Vickers [1988], the strategic delegation literature has shown how vertical restraints, strategic pricing, or specialized agency contracts can be used to distort firm behavior in a profitable way.[9] Section 5 synthesizes this literature using the notion of conjectural consistency. We develop a very general model of "distorted" games in which upstream firms can strategically distort the behavior of downstream firms. The model nests essentially all possible mechanisms for strategic distortion, including strategic delegation, strategic double marginalization, and vertical restraints like resale price maintenance. The optimal way to distort a downstream firm is always to make it behave *as if* it has a consistent conjecture, no matter the nature of the distortion. Consistent with the insights from the strategic delegation literature, any kind of strategic distortion of a downstream firm's behavior will induce exactly the same downstream equilibrium actions (holding the rival's reaction behavior constant), namely the distorted firm's preferred profile along the rival's reaction curve.

If upstream competitors strategically distort the downstream firms in parallel, they induce both downstream firms to act as if they have (globally) consistent conjectures, at least in equilibrium. This induces them play the equilibrium strategies corresponding to a consistent conjectures equilibrium (CCE) – a concept introduced by Bresnahan [1981] – or else a close analogue which we call an "induced CCE."[10] This is a surprising result, since the CCE derives from non-Nash behavior,[11] leading many to criticize the concept. But we show that CCEs arise in games with strategically distorted behavior even though every firm engages in ordinary Nash behavior.

[9] See also Gal-Or [1991]; Rey and Stiglitz [1995]; O'Brien and Shaffer [1992]; Mathewson and Winter [1984]; Jansen et al. [2007].

[10] Specifically, upstream competition through strategic distortions always generates an induced CCE, which is an equilibrium such that each firm is distorted into playing the same equilibrium strategy it would adopt if it had a (globally) consistent conjecture, given its rival's (distorted) reaction behavior. This may or may not generate exactly the same equilibrium strategies as Bresnahan's "perfect" CCE. But it always has the property that each firm is properly accounting for strategic effects in equilibrium, making it a close analogue.

[11] To our knowledge, the only conventional Nash equilibrium in a supermodular game (without strategic distortions) that happens to be a CCE is the Bertrand equilibrium with undifferentiated products. See Bresnahan [1981].

Section 6 concludes by discussing implications for antitrust, vertical integration, and the theory of the firm. That joint profits can be credibly enhanced because of indirect strategic effects may imply a role for a "strategic theory of the firm." In most situations when one firm imposes a positive or negative externality on another firm, that effect can be internalized either by Coasean bargaining or by integration. However, when vertically related firms integrate to align interests, third party rivals will understand that the downstream firm's equilibrium behavior, previously distorted by noncooperative double marginalization, will be replaced by the integrated firm's more competitive behavior, and this may provide lower joint profits. Thus, counterintuitively, joint profits may be higher when the firms are separate and noncooperative than when they can bargain efficiently.

As some portions of this article address the theory of double marginalization at an abstract level, the reader primarily interested in antitrust implications will find Sections 2, 3, and 6 most relevant.

2 Motivation and Examples

We begin with three examples illustrating profitable double marginalization within familiar competitive environments. In each example, we modify a downstream duopoly game so that firm 1 must pay a linear per-unit "transfer price" of t to an upstream firm. First, the upstream firm chooses t noncooperatively – meaning that it can charge only the distortionary linear price t, not a two-part tariff – and then the downstream firms compete in simultaneous moves, with firm 1 treating t as a marginal cost. Upstream costs are normalized to zero, so $t = 0$ corresponds to the absence of double marginalization and, by extension, vertical integration. This is our basis for comparing joint profits (of firm 1 and the upstream firm). The examples are arranged in increasing order of the joint profit increase. We begin with a simple game of Hotelling price competition.

Example 1 (Hotelling): There is a unit line of locations $x \in [0, 1]$, with one consumer at each location, and with firms 1 and 2 located at $x = 0$ and $x = 1$, respectively. Each $i = 1, 2$ chooses a price, p_i. A consumer at location x gets surplus $v - p_1 - x$ if it buys from firm 1 and $v - p_2 - (1 - x)$ if it buys from firm 2, where $v \geq 3$.[12] This yields demand functions $q_i = \frac{1}{2}[1 - p_i + p_j]$ for each i and $j \neq i$. Firm 2's costs are zero; firm 1 must pay a per-unit transfer price t to the upstream firm. Best response functions are thus $R_1(p_2|t) = \frac{1}{2}(1 + t + p_2)$ and $R_2(p_1) = \frac{1}{2}(1 + p_1)$. Given t, firm 1's equilibrium price and output are

$$p_1^*(t) = \frac{3 + 2t}{3} \qquad q_1^*(t) = \frac{3 - t}{6} \tag{1}$$

Since t is just a transfer, joint profits of firm 1 and the upstream firm are $J(t) = q_1^*(t)p_1^*(t)$. As is well-established in the strategic delegation literature, J is maximized with the transfer price t^S that induces $p_1^*(t^S) = p_1^S$, where $p_1^S \equiv \frac{3}{2}$ is firm 1's Stackelberg price when $t = 0$. The upstream firm

[12] The constraint on v simply ensures that every consumer will always get nonnegative surplus from each product in equilibrium.

prices noncooperatively, maximizing only upstream profits, $tq_1^*(t)$, resulting in a monopoly price t^m. This price will enhance joint profits if and only if it lies in the "Goldilocks zone" – the interval of transfer prices $[0, \bar{t}]$ that provide weakly larger joint profits than integration ($t = 0$). Here $\bar{t}$ is defined implicitly by $J(\bar{t}) = J(0)$. It is easy to verify that

$$t^m = \bar{t} = \frac{3}{2} \qquad t^S = \frac{3}{4} \tag{2}$$

Thus double marginalization with an upstream monopoly price provides exactly the same joint profits as vertical integration.[13] Hence noncooperative vertical pricing will *never* reduce joint profits in a simple Hotelling environment, even if the upstream firm is a monopolist. Further, if the upstream firm is price constrained to any degree, no matter how tiny – for instance, there may be a third party selling the same input – double marginalization will strictly raise joint profits. Figure 1 provides a graph illustrating the Goldilocks zone. It shows the best response functions associated with $t = 0$ and $t = \bar{t} = t^m$. Joint profits weakly rise for any upstream markup t that leaves firm 1's distorted best response function somewhere in the middle. The dotted line IP_1 is the iso-profit curve that runs through these equilibria, while SE denotes the Stackelberg equilibrium at which joint profits are maximized. ■

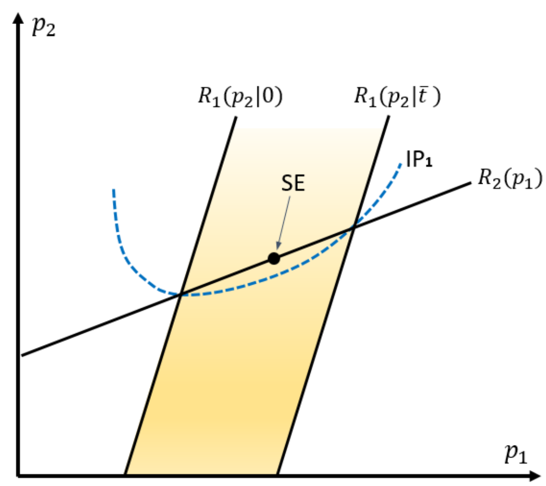

Figure 1: Goldilocks Zone

If the model involved linear Bertrand competition with differentiated products (and the same costs), the analysis is very similar to Example 1, except that double marginalization is weakly profitable only if the upstream markup is strictly price-constrained.[14] Importantly, however, if the model were instead linear Cournot, any positive transfer price – and hence any level of double

[13] The result would hold up if we changed the length of the Hotelling line or the marginal "travel cost" (which is presently 1). It similarly holds up if we allow for asymmetries in v or if we give the firms asymmetric internal production costs.

[14] Specifically, if the demand specification were $q_i = 1 - p_i + sp_j$ for each i, where $s \in (0, 1)$, then we would obtain the result $0 < \bar{t} < t^m$, and thus double marginalization will strictly increase joint profits only if the upstream firm is sufficiently price-constrained.

marginalization – would reduce joint profits. Intuitively, a linear Cournot firm wants to commit to a *more* competitive strategy (a lower output level), while a Bertrand or Hotelling firm wants to commit to a *less* competitive strategy (a higher price). Since double marginalization always induces less competitive behavior, it has opposite profit effects in these alternative cases.

Strictly profitable double marginalization does not always require that the upstream firm is price-constrained, however. Example 2 demonstrates this using an iso-elastic Cournot game. Here even unconstrained double marginalization can strictly increase joint profits.

Example 2 (Iso-Elastic Cournot): Each firm $i = 1, 2$ chooses an output level $q_i \geq 0$. Products are undifferentiated, and aggregate inverse demand takes the constant elasticity form $p(Q) = 1/Q$, where $Q = q_1 + q_2$ is aggregate output.[15] Marginal costs are constant and equal to c_2 for firm 2 and $c_1 + t$ for firm 1, where $c_1, c_2 > 0$. Thus payoffs are $\pi_1 = q_1 p(q_1 + q_2) - (c_1 + t)q_1$ and $\pi_2 = q_2 p(q_1 + q_2) - c_2 q_2$ for firms 1 and 2, respectively. These yield the best response functions

$$R_1(q_2|t) = \sqrt{\frac{q_2}{c_1 + t}} - q_2 \qquad R_2(q_1) = \sqrt{\frac{q_1}{c_2}} - q_1 \tag{3}$$

Unlike most Cournot models, these best response functions are not monotonically decreasing, but rather take an inverted-U shape. Conditional on t, the Nash equilibrium output levels and price are

$$q_1^*(t) = \frac{c_2}{(c_1 + t + c_2)^2} \qquad q_2^*(t) = \frac{c_1 + t}{(c_1 + t + c_2)^2} \qquad p^*(t) = c_1 + t + c_2 \tag{4}$$

and equilibrium profits for the downstream firms are $\pi_1^*(t) = c_2 q_1^*(t)$ and $\pi_2^*(t) = (c_1 + t)q_2^*(t)$. Joint profits are $J(t) \equiv \pi_1^*(t) + t q_1^*(t)$, maximized with the transfer price t^S that induces firm 1's Stackelberg output: $q_1^*(t^S) = q_1^S \equiv c_2/4c_1^2$. Note that $q_1^S < q_1^*(0)$ if and only if $c_1 > c_2$. Thus, in contrast to ordinary Cournot competition, a firm may wish to commit to a less competitive strategy, provided it is the less efficient producer. In such a case, the highest positive transfer price that weakly enhances joint profits is $\bar{t}$, defined implicitly by $J(\bar{t}) = J(0)$. As in the previous example, let t^m denote the monopoly transfer price, which maximizes $t q_1^*(t)$. It is easy to verify that the transfer prices $(t^S, t^m, \bar{t})$ are unique and given by

$$t^S = c_1 - c_2 \qquad t^m = c_1 + c_2 \qquad \bar{t} = \frac{c_1^2 - c_2^2}{c_2} \tag{5}$$

Unlike the linear Bertrand and Cournot models, where the jointly-optimal markup is either always positive or always negative, here it can be either positive, negative, or zero. Note that that $t^S < t^m$ for all possible cost levels, so double marginalization never maximizes joint profits.

[15]This demand specification has been studied in a number of theoretical papers on duopoly pricing. See, e.g., Puu (1991). Note that, while this model is problematic in the monopoly case (a monopolist would set an infinite price), it creates no such problems in the duopoly case, provided both firms have positive production costs.

Nevertheless, t^m lies in the Goldilocks zone $[0, \bar{t}]$ whenever $c_1 \geq 2c_2$.[16] ■

Example 2 shows that in some games, unconstrained double marginalization strictly improves, though does not maximize, joint profits. Even this modest limitation does not always arise. Example 3 illustrates a game in which the monopoly transfer price may inadvertently maximize joint profits, or even be too *low* to achieve the joint-optimum.

Example 3 (Freeriding Innovators): There are two firms engaging in costly research to complete a new invention worth payoff 1. Each firm chooses a costly research success probability $\phi_i \in [0, 1]$, which gives the probability that its own research will solve its own problem, at convex cost $\beta\phi_i^2$, where $\beta \geq 1$. In addition to this internal research cost, Firm 1 must make a linear transfer payment $t\phi_1$ to an upstream firm. This transfer represents, for example, that firm 1's research requires use of a patented process for which an upstream patent holder charges a royalty of t.

There are information spillovers, however, and thus one firm's research may inadvertently solve the other's problem. In particular, firm j's research will solve problem i (independently of whether it solves problem j) with probability ϕ_j, and hence problem i is solved with probability $\phi_i + (1-\phi_i)\phi_j$. Therefore payoff functions are $\pi_1 = \phi_1 + (1-\phi_1)\phi_2 - \beta\phi_1^2 - t\phi_1$ and $\pi_2 = \phi_2 + (1-\phi_2)\phi_1 - \beta\phi_2^2$ for firms 1 and 2, respectively. The corresponding best response functions are $R_1(\phi_2|t) = (1-t-\phi_2)/2\beta$ and $R_2(\phi_1) = (1-\phi_1)/2\beta$. Conditional on t, the Nash equilibrium is

$$\phi_1^*(t) = \frac{2\beta - 1 - 2\beta t}{4\beta^2 - 1} \qquad \phi_2^*(t) = \frac{2\beta - 1 + t}{4\beta^2 - 1} \tag{6}$$

and profits are $\pi_i^*(t) = \beta\phi_i^*(t)^2 + \phi_j^*(t)$ for each downstream firm i. Joint profits of firm 1 and the upstream firm are maximized with the transfer price t^S that induces firm 1's Stackelberg strategy ϕ_1^S, where

$$\phi_1^S = \frac{\beta - 1}{2\beta^2 - 1} \qquad t^S = \frac{2\beta - 1}{4\beta^2 - 2} \tag{7}$$

The upstream monopoly transfer price which maximizes $t\phi_1^*(t)$ is $t^m = (2\beta - 1)/4\beta$. It is easy to verify that $sign\{t^S - t^m\} = sign\{\widetilde{\beta} - \beta\}$, where $\widetilde{\beta} \equiv (1+\sqrt{3})/2 \simeq 1.366$. Hence, if $\beta = \widetilde{\beta}$, then $t^m = t^S$ and thus the monopoly transfer price inadvertently maximizes joint profits. If $\beta \in [1, \widetilde{\beta})$, then t^m is too *low* to maximize joint profits, a possibility that never arises in any of the preceding examples. The Goldilocks zone in this game is $[0, \bar{t}]$, where $\bar{t} = (2\beta - 1)/(2\beta^2 - 1)$. It is then straightforward to verify that $t^m \in [0, \bar{t}]$ if and only if $\beta \leq \overline{\beta} \simeq 2.22$.

[16]Due to the model's exotic best response functions, its equilibrium becomes unstable when costs are highly asymmetric. Fortunately, all of the above results can still obtain in stable equilibria. To show this, we make use of Puu [1991], which studied precisely the same game (albeit without double marginalization), and showed that its equilibrium is stable if and only if $\bar{c}/\underline{c} < 3 + \sqrt{2} \simeq 5.82$, where $\bar{c}$ ($\underline{c}$) is the maximum (minimum) of c_1 and c_2. With this, assume $c_1 = (2+\alpha)c_2$ for some $\alpha > 0$, implying $t^m < \bar{t}$. Then the equilibrium with $t = 0$ is stable if $c_1/c_2 = 2 + \alpha \leq 5.82$, which is true for small α. Similarly, the equilibrium with $t = t^m$ is stable if $(c_1 + t^m)/c_2 = 5 + 2\alpha \leq 5.82$, and this too is satisfied for small α.

To explain these results, note that $\phi_1^S < \phi_1^*(0)$, reflecting that the counterfactual integrated firm would like to commit to a freeriding strategy in which it lets firm 2 do most of the work. In fact, $\phi_1^S = 0$ when $\beta = 1$, in which case firm 1 would like to commit to sitting out entirely. Of course, an upstream firm that prices noncooperatively would never want to induce this outcome, since it generates an upstream profit of zero. Thus, for β close to 1, the jointly-optimal strategy is *less* competitive than that induced by double marginalization. The opposite is true when β is sufficiently high, and hence continuity ensures there is some intermediate β-value at which the monopoly transfer price inadvertently maximizes joint profits. ■

These examples, in addition to the briefly discussed cases of differentiated Bertrand and linear Cournot, show that the possible profit effects of double marginalization run the gamut: it can be categorically harmful, joint profit-maximizing, or anything in between. Examples 2 and 3 further show that this distinction does not simply hinge on the distinction between strategic complements and strategic substitutes. Example 2 involves neither strategic complements nor strategic substitutes, since reaction functions are non-monotonic. Further, the best response functions in Example 3 are identical to those arising in linear Cournot competition, notwithstanding that these models engender completely opposite conclusions about double marginalization. We return to this point in Section 3.

2.1 A Note on Horizontal Double Marginalization

We now briefly show that our results carry over to the case of "horizontal" double marginalization, which arises when horizontally related firms set independent markups for goods that are complementary in consumption.[17] Horizontal double marginalization is generally studied under the assumption that firms face no third party competition (we are unaware of a counterexample.) But we show in Example 4 that under third party competition, horizontal (and noncooperative) double marginalization can raise profits. Importantly, this demonstrates that our basic results are not driven by an artificial change in the timing of the game.

Example 4 (Competing Bundles): Let there be two bundles, A and B, which are imperfect substitutes that compete in prices. Each bundle contains a pair of perfectly complementary goods. The total price of bundle $I \in \{A, B\}$ is denoted p_I. Demand for bundle I is $q_I = 1 - p_I + sp_J$, where $J \neq I$ and $s \in (0, 1)$. Let $\Pi_I = q_I p_I$ denote the total profits of bundle I. Bundle B consists of two goods produced by a single integrated firm, which therefore chooses a single price p_B to maximize Π_B. The components of bundle A are priced noncooperatively by two firms $i = 1, 2$. Thus $p_A = \lambda_1 + \lambda_2$, where λ_i is the component price that firm i sets to maximize its own profits,

[17] As a game-theoretic matter, the difference between horizontal and vertical double marginalization is that the complementary goods are priced simultaneously rather than sequentially.

$\lambda_i q_A$, taking (λ_j, p_B) as given. All prices are set simultaneously. The equilibrium involves

$$\lambda^* = \frac{2+s}{2(3-s^2)} \qquad \Pi_A^* = 2(\lambda^*)^2 \tag{8}$$

Now consider the alternative equilibrium arising when bundle A is priced by an integrated firm, eliminating horizontal double marginalization. In this case the equilibrium has

$$p_A^{**} = \frac{1}{2-s} \qquad \Pi_A^{**} = (p_A^{**})^2 \tag{9}$$

Thus $p_A^* \equiv 2\lambda^* > p_A^{**}$, reflecting that horizontal double marginalization makes bundle A more expensive. However, we also find that $\Pi_A^* > \Pi_A^{**}$ whenever $s > \underline{s}$, where $\underline{s} = \sqrt{2-\sqrt{2}} \simeq .76$. Thus, horizontal double marginalization strictly increases joint profits earned on bundle A when the competing bundles are sufficiently close substitutes. ■

3 General Model

We now present a generalized double marginalization game. There is a downstream market with two firms, $i = 1, 2$.[18] There is also an upstream firm that double marginalizes firm 1. These three firms are engaged in a two-stage game, which is a generalization of the format used in Examples 1-3. The stages in the generalized game are defined as follows.

Stage 2 (Downstream Competition)

In stage 2, the downstream firms simultaneously choose strategies $a_1 \in A_1, a_2 \in A_2$ to maximize payoff functions $\pi_1(a_1, a_2) - tX_1(a_1, a_2)$ and $\pi_2(a_2, a_1)$. Here $tX_1(a_1, a_2)$ gives firm 1's total transfer payment as a function of strategies and a *transfer price* t chosen by the upstream firm.[19] The function $X_1 : A_1 \times A_2 \to \mathbb{R}_+$ is the *transfer base*, which we can interpret either as firm 1's output ($X_1 = q_1(a_1, a_2)$), or more generally as some strictly increasing transformation of firm 1's output. The transfer price satisfies $t \leq \tau$ for some $\tau > 0$, which allows for the possibility that the upstream firm is price-constrained. Best response functions are denoted $R_1(a_2|t)$ and $R_2(a_1)$ for firms 1 and 2, respectively. We make the following assumptions:

(A1): each A_i is a compact interval in $\mathbb{R}$

(A2): each π_i is twice continuously differentiable in (a_i, a_j) and strictly concave in a_i

(A3): X_1 is twice continuously differentiable in (a_1, a_2) and strictly monotonic in a_1

[18] The results that follow can be extended in a straightforward way to aggregate games with $N > 2$ downstream firms. The duopoly assumption is used for expositional simplicity.

[19] X_1 is written as a function of both a_1 and a_2 because in some games, e.g. Bertrand competition, output depends on both players' strategies.

(A4): for every $t \in [0, \tau]$, there exists a unique and interior downstream equilibrium $(a_1^e(t), a_2^e(t))$

(A5): The Stackelberg objective function $\pi_1(a_1, R_2(a_1))$ is strictly concave with a unique and interior maximizer a_1^S, and there exists some $t^S \in \mathbb{R}$ such that $a_1^e(t^S) = a_1^S$

Assumptions (A1) and (A2) ensure that the best response functions are well-defined and continuously differentiable. Assumption (A3) ensures that t acts like a marginal cost, so that $R_1(\cdot|t)$ is shifting strictly monotonically in t.[20] Since R_2 is independent of t, this implies that $a_1^e(t)$ is strictly monotonic in t. Assumption (A4) ensures that equilibria are unique and pinned down by first order conditions. Assumptions (A5) ensures the Stackelberg equilibrium is unique and interior, and that there is some t^S that would induce this equilibrium. Note, however, that we make no assumption about the sign or magnitude of t^S. Finally, the concavity provision in (A5) simplifies the comparative statics by ensuring that joint profits are monotonically falling as t moves further away from t^S.

Stage 1 (Noncooperative Upstream Pricing)

In stage 1, the upstream firm sells a costless input to firm 1 according to the transfer price schedule tX_1. It sets the transfer price t to maximize equilibrium transfer revenue $T(t) \equiv tX_1(a_1^e(t), a_2^e(t))$ under the constraint $t \leq \tau$. The transfer base function X_1 and the maximal transfer price $\tau > 0$ are exogenously given. We make one additional assumption ensuring that the upstream firm's maximization problem is well-behaved:

(A6) $T(t)$ is strictly concave with a unique unconstrained maximizer $t^m \in (0, \infty)$.

Let $\widehat{t} = \min\{t^m, \tau\}$ denote the transfer price that maximizes transfer revenues subject to the constraint $t \leq \tau$. We refer to the equilibrium with $t = 0$ as the *integration equilibrium*, reflecting that a vertically integrated firm would ignore t.

Let $\pi_1^S(a_1) \equiv \pi_1(a_1, R_2(a_1))$ give firm 1's profit function when it acts as a Stackelberg leader with $t = 0$. Since T is just a transfer, equilibrium joint profits are $\pi_1(a_1^e(t), a_2^e(t)) = \pi_1^S(a_1^e(t))$. That is, joint profits of the upstream firm and firm 1 are equivalent to the profits an integrated firm would get if it were committed to playing $a_1^e(t)$. In the following decomposition, we show conditions under which *some* positive degree of double marginalization is jointly-profitable, implying at the very least that equilibrium joint profits will rise if the upstream firm is sufficiently price-constrained.

Proposition 1. *There exists $\tau > 0$ such that double marginalization strictly increases joint profits if and only if*

$$\frac{\partial \pi_1}{\partial a_2} \times \frac{\partial R_2}{\partial a_1} \times \frac{\partial a_1^e}{\partial t} > 0$$

when evaluated at the integration equilibrium.

[20] If $\partial X_1/\partial a_i > (<) 0$, then R_1 shifts downward (upward) in t, so we can interpret a_i as output (price).

Proof: The total derivative of joint profits with respect to t is

$$\frac{d\pi_1^S(a_1^e(t))}{dt} = \frac{\partial \pi_1(a_1^e, a_2^e)}{\partial a_1}\frac{\partial a_1^e}{\partial t} + \frac{\partial \pi_1(a_1^e, a_2^e)}{\partial a_2}\frac{\partial R_2(a_1^e)}{\partial a_1}\frac{\partial a_1^e}{\partial t} \tag{10}$$

where we have omitted the argument t on the righthand side. When $t = 0$, firm 1's first order condition requires that $\frac{\partial \pi_1}{\partial a_1} = 0$. Hence, locally at the integration equilibrium, joint profits are strictly increasing in t if $\frac{\partial \pi_1}{\partial a_2}\frac{\partial R_2}{\partial a_1}\frac{\partial a_1^e}{\partial t} > 0$. Thus, if τ is sufficiently small then $\widehat{t}$ will strictly increase joint profits. For the "only if," note that $d\pi_1^S(a_1^e(t))/dt < 0$, since π_1^S is strictly concave in a_1 and a_1^e is strictly monotonic in t. Thus, if the stated inequality were not true, then any positive t would erode joint profits. *Q.E.D.*

Proposition 1 shows that the profitability of noncooperative double marginalization depends on the interaction of three local conditions: strategic complementarity ($\frac{\partial R_2}{\partial a_1}$); payoff complementarity ($\frac{\partial \pi_1}{\partial a_2}$); and the direct impact of raising firm 1's costs ($\frac{\partial a_1^e}{\partial t}$). When either one or three of these things are positive (and none are zero), some degree of double marginalization is beneficial. For example, linear Cournot is a game of strategic substitutes and payoff substitutes where higher costs lead to lower quantity (all three partial derivatives are negative), whereas the Freeriding Innovators games is one of strategic substitutes and payoff complements where higher costs lead to less innovation (one positive derivative, two negative ones).

The previous proposition shows when there exists some transfer price such that profits improve under double marginalization, but when does fully unconstrained double marginalization improve profits? The relative values of the unconstrained monopoly transfer t^m and the jointly-optimal transfer t^S can be investigated by decomposing marginal joint profits with respect to t – given by $\frac{d}{dt}\pi_1(a_1^e(t), a_2^e(t)) = \frac{d}{dt}\pi_1^S(a_1^e(t))$ – into two distinct marginal effects of changing t. The first is the *direct effect*, $DE(t)$, that derives from the resulting marginal change in $a_1^e(t)$. The second is the the *strategic effect*, $SE(t)$, that derives from the corresponding marginal change in $a_2^e(t) = R_2(a_1^e(t))$. Explicitly:

$$DE(t) = \frac{\partial \pi_1(a_1^e(t), a_2^e(t))}{\partial a_1} \times \frac{\partial a_1^e(t)}{\partial t} \tag{11}$$

$$SE(t) = \frac{\partial \pi_1(a_1^e(t), a_2^e(t))}{\partial a_2} \times \frac{\partial R_2(a_1^e(t))}{\partial a_1} \times \frac{\partial a_1^e(t)}{\partial t} \tag{12}$$

Note that $DE(0) = 0$, which is an implication of firm 1's first order condition in the integration equilibrium, and $DE(t) < 0$ for all $t > 0$ (and vice versa). To see this, fix $t > 0$ and note that $\partial X_1/\partial a_1 > (<)0$ implies $\partial R_1/\partial t < (>)0$ and by extension $\partial a_1^e/\partial t < (>)0$. On the other hand, the strategic effect $SE(0)$ could be positive, negative, or zero. The sign of $SE(0)$ equals that of t^S, which is a direct implication of Proposition 1 (note that the inequality in Proposition 1 simply says $SE(0) > 0$). Thus the sign of $SE(0)$ is negative when double marginalization is categorically harmful and positive when at least some degree of double marginalization is beneficial. The case

$SE(0) = 0$ arises when firm 1 cannot benefit from commitment (e.g., when firm 1 is a monopolist rather than a duopolist, hence $\partial\pi_1/\partial a_2 = 0$). The following proposition establishes that double marginalization via the unconstrained monopoly transfer t^m raises joint profits when beneficial strategic effects are strong enough to counteract negative direct effects.

Proposition 2. *Unconstrained double marginalization enhances joint profits if and only if the strategic effect is sufficiently large compared to the direct effect. In particular,*

1) $sign\{t^S - t^m\} = sign\{SE(t^m) + DE(t^m)\}$.

2) If $SE(\cdot) + DE(\cdot)$ shifts vertically upward, t^S rises in relation to t^m.

Proof: The first order conditions that pin down t^S and t^m are, respectively:

$$SE(t^S) + DE(t^S) = 0 \qquad T'(t^m) = 0 \tag{13}$$

$SE(t) + DE(t) = d\pi_1^S(a_1^e(t))/dt$ is strictly decreasing (by strict concavity of π_1^S in a_1 and monotonicity of $a_1^e(t)$ in t), and $T'(t)$ is strictly decreasing (by strict concavity of T). Thus $t^m > (\leq) t^S$ if and only if $SE(t^m) + DE(t^m) \leq (\geq) 0$. Holding $T'(\cdot)$ constant, if $SE(\cdot) + DE(\cdot)$ shifts upward, so that $SE(t^m) + DE(t^m)$ is higher, then t^S must rise in relation to t^m. *Q.E.D.*

Let $\pi_1^D(t) = \pi_1^S(a_1^e(t)) - T(t)$ denote the downstream component of joint profits, given t. Then since $T(.)$ is maximized at t^m by definition, $\partial\pi_1^D(t^m)/\partial t = SE(t^m) + DE(t^m)$. Thus, when $t^m = t^S$, downstream marginal profits (with respect to t) are zero in equilibrium. If $t^m < t^S$, then downstream profits are *increasing* in t within a neighborhood of t^m – an unusual result that reflects a very strong strategic effect. For instance, this is true in Example 3 when $\beta \in [1, \widetilde{\beta})$. Even if $t^m > t^S$, strategic effects can be sufficiently strong that joint profits are higher at t^m than at $t = 0$, as we saw in the examples in Section 2. The condition for joint profits to improve is weaker than the condition that joint profits are maximized, namely requiring that t^m lies in the interval $[0, \bar{t}]$, where $\bar{t} = \max\{t \geq 0 | \pi_1^S(a_1^e(t)) = \pi_1^S(a_1^e(0))\}$. For example, if $t^S > 0$ and $SE(t) + DE(t)$ is linear in t, then this simply requires that $t^m \leq 2t^S$.

Simple price theory can show that t^S and t^m may be arbitrarily different no matter what the equilibrium or best response functions are in the downstream game. Therefore, via Proposition 2, the relation between direct effects and strategic effects is similarly ambiguous, and hence the profitability of double marginalization depends on more than just characteristics of the downstream game. To see this, note that the monopoly transfer price t^m is determined solely by the elasticity of demand for the input good. Letting $X_1(t) \equiv X_1(a_1^e(t), a_2^e(t))$, the elasticity of demand for the upstream good is

$$\epsilon_{X_1}(t) = X_1'(t)\frac{t}{X_1(t)} \tag{14}$$

and hence when the optimal transfer price is interior, t^m is chosen such that

$$T'(t) = X_1(t)\big[1 + \epsilon_{X_1}(t)\big] = 0 \tag{15}$$

Therefore, the well-known inverse elasticity rule holds: transfer revenue is maximized if and only if $\epsilon_{X_1}(t) = -1$. The elasticity $\epsilon_{X_1}(t)$, and hence t^m, depends solely on equilibrium strategies conditional on t: for any two games where equilibrium strategies as a function of t are identical, elasticities are likewise identical for all t, and hence monopoly transfer prices t^m are identical. The joint profit maximizing transfer t^S, however, may vary arbitrarily based on differences in strategic effects. As such, knowing only properties relevant to the best response functions or equilibrium actions downstream, such as whether the game is one of strategic substitutes or complements, is not enough to know whether double marginalization is profitable or not. Further, two games may appear very similar in certain respects, and yet double marginalization can have very different profit effects between them. To illustrate, consider the Freeriding Innovators game (Example 3) with $\beta = 1$. Firm 1's equilibrium input demand, $\phi_1^*(t) = (1 - 2t)/2$, is identical to that arising in a Cournot duopoly with inverse demand $p(q_1 + q_2) = 1 - q_1 - q_2$ and constant marginal costs of t for firm 1 and zero for firm 2. Indeed, best response functions for both firms are identical in these games. Despite this similarity, double marginalization has totally different profit effects in these settings, with $t^S < 0 = \bar{t} < t^m$ in the Cournot game and $0 < t^m < t^S < \bar{t}$ in the Freeriding game. Proposition 3 provides a more formal expression.

Proposition 3. *Suppose the functions (π_1, π_2) are changed to some alternatives $(\widetilde{\pi}_1, \widetilde{\pi}_2)$ such that, for every t, the resulting downstream equilibrium strategies are unchanged. Let $\widetilde{t}^m$ and $\widetilde{t}^S$ denote the values of t^m and t^S in this alternative game. Then $\widetilde{t}^m = t^m$, but $\widetilde{t}^S$ and t^S may be arbitrarily different. The same is true if we further require that best response functions are unchanged.*

Proof: The upstream firm maximizes $tX_1(a_1^e(t), a_2^e(t))$. Equilibrium strategies as a function of t are unchanged by hypothesis, and thus $\widetilde{t}^m = t^m$. However, cross-derivatives $\partial\widetilde{\pi}_i/\partial a_j$ (and potentially reaction functions) can differ in the alternative game even if equilibrium actions conditional on t are unchanged. For example, let $\widetilde{\pi}_i = \pi_i + f_i(a_j)$. Since $\partial\widetilde{\pi}_i/\partial a_i = \partial\pi_i/\partial a_i$, direct effects are identical under either payoff function, but strategic effects – and hence the jointly-optimal transfer price – can be arbitrarily different depending on the magnitude of the derivatives $f_i'(a_j)$. *Q.E.D.*

4 Insights from Conjectural Variation

In Section 3, we showed that the profitability of noncooperative double marginalization is in general ambiguous. For example, the jointly-optimal upstream markup is negative if competition is linear Cournot, but positive if competition is differentiated Bertrand. On the other hand, if competition is iso-elastic Cournot, then the optimal upstream markup could be positive, negative, or zero. We showed that the potential profitability of double marginalization depend on three separate partial derivatives, but this section seeks to provide a richer explanation. We achieve this using insights from the conjectural variation (CV) literature, and in particular the notion of conjectural consistency.[21]

First we provide a very short overview of CV. A firm's "conjecture" is its implicit belief about the rate at which its rival will respond to changes in its own conduct. Conjectures are traditionally defined with respect to output, and expressed as a derivative of rival-output with respect to own-output. Adhering to this convention, a conjecture for firm $i \in \{1, 2\}$ is some (possibly nonconstant) term $\nu_i = \partial q_j / \partial q_i$, where q_j is rival output and q_i is i's own output. To illustrate, contrast the Cournot and Bertrand models. A Cournot firm has a degenerate conjecture ($\nu_i = 0$), because it takes rival output as fixed when fashioning a best response. But a Bertrand firm believes that if it increases its own output (by lowering its price), then it will steal some sales from its rival ($\nu_i < 0$).

A CV game is a standard simultaneous-move game that has been extended so that firms' conjectures are explicit parameters that can be altered without changing the rules of the game. Consider the following CV game with differentiated duopoly quantity competition. Each firm i chooses an output level q_i. Inverse demand for firm i is $p_i = 1 - q_i - sq_j$, where $s \in (0, 1)$ denotes the degree of product substitutability. Production costs are normalized to zero, so i's payoff is $\pi_i(q_i, q_j) = (1 - q_i - sq_j)q_i$. Each firm i is assumed to have a (constant) conjecture $\nu_i = \partial q_j / \partial q_i$. This enters into the firm's optimization problem, with each i's first order condition given by

$$\frac{d\pi_i(q_i, q_j)}{dq_i} = \frac{\partial \pi_i(q_i, q_j)}{\partial q_i} + \frac{\partial \pi_i(q_i, q_j)}{\partial q_j} \nu_i = 1 - 2q_i - sq_j - s\nu_i q_i = 0 \tag{16}$$

Thus i's response function is given by $\widetilde{R}_i(q_j | \nu_i) = (1 - sq_j)/(2 + s\nu_i)$. We refer to an intersection of these response functions as a *conjectural variations equilibrium* (CVE). If $\nu = \nu^C = 0$, the CVE is Cournot-Nash; if $\nu^B = -s$, the CVE is the Bertrand equilibrium.[22]

Following Bresnahan [1981], a conjecture is *consistent* if it is correct, meaning that it equals the slope of the rival's best response function. Letting $\widetilde{r}(\nu_i) \equiv -s/(2 + s\nu_i)$ denote the (constant) slope of $\widetilde{R}_i$ conditional on ν_i, this means that ν_i is consistent if $\nu_i = \widetilde{r}(\nu_j)$. If only one firm has a consistent conjecture, the result is a Stackelberg equilibrium.[23] If the firms' conjectures are mutually-consistent, then the resulting equilibrium is a *consistent conjectures equilibrium* (CCE). (Technically the CCE concept in Bresnahan [1981] requires only that both conjectures are consistent within an open neighborhood of the equilibrium, but this is equivalent to global consistency in a linear game with constant conjectures.) By symmetry, a CCE in this linear model involves the self-consistent conjecture $\nu^\star = \widetilde{r}(\nu^\star)$. Solving this quadratic yields $\nu^\star = -(1 - \sqrt{1 - s^2})/s$. Since $s \in (0, 1)$, we obtain the strict ordering

[21] The theory of conjectural variations was first formulated by Bowley [1924]. Bresnahan [1981] was among the first to study models with consistent conjectures.

[22] To see this, note that inverting the demand system (p_1, p_2) yields Bertrand-style demand functions $q_i = \gamma[1 - s - p_i + sp_j]$ for each i, where $\gamma \equiv (1 - s^2)^{-1}$. Thus, in the Bertrand game, firm i it believes that if it increases q_i by 1 (by cutting p_i by γ^{-1}), the resulting change in q_j will be $\gamma[s(-\gamma^{-1})] = -s$.

[23] For example, the Stackelberg-Cournot equilibrium arises when one firm (the follower) has conjecture ν^C and the other firm (the leader) has conjecture $\widetilde{r}(\nu^C)$.

$$-1 < \nu^B < \widetilde{r}(\nu^B) < \nu^\star < \widetilde{r}(\nu^C) < \nu^C = 0$$

and hence both Cournot and Bertrand competition involve inconsistent conjectures.[24] In general, CCEs (and CVEs with arbitrary conjectures) are not Nash equilibria of the underlying game. The artificiality of conjectures is a principal reason why many economists have criticized the technique.[25] However, we show that conjectural variation can be useful for interpreting Nash equilibria, and can serve as a reduced form for analyzing games with strategically distorted behavior.

A firm with a consistent conjecture naturally accounts for strategic effects, which is valuable to the firm. If its conjecture is inconsistent, then it will either overstate or understate its rival's competitiveness, as captured by the slope of the rival's reaction function. To illustrate, note that a higher (lower) conjecture about rival output means the firm believes its rival will be more accommodating (competitive). Thus if the firm's conjecture is inconsistently low (high), then it is overstating (understating) rival competitiveness, and can therefore benefit from being induced to behave less (more) competitively.

This is what explains why the jointly-optimal upstream markup looks so different within different games. In linear Cournot, the firm's conjecture is inconsistently high ($\nu^C > \widetilde{r}(\nu^C)$), suggesting it is understating competition, and can benefit from behaving more competitively. Thus, the jointly-optimal upstream markup is negative ($t^S < 0$). By contrast, the Bertrand firm's conjecture is inconsistently low ($\nu^B < \widetilde{r}(\nu^B)$), implying it is overstating rivals' competitiveness and can thus benefit from behaving less competitively, which can be accomplished by applying a positive upstream markup ($t^S > 0$). In either case, the jointly-optimal markup leads the downstream firm to play the same equilibrium strategy it would adopt if it had a consistent conjecture. An interpretation is that the objective of strategic delegation is, in effect, to correct an inconsistent conjecture. Noncooperative double marginalization enhances joint profits when it inadvertently leads the downstream firm to play as if it has a "better" conjecture than its actual (degenerate) conjecture.

It is easy to see all of these possibilities within a single framework by returning to the isoelastic Cournot game from Example 2. Recall that the jointly optimal transfer price is $t^S = c_1 - c_2$, and thus the sign of this price depends on the comparison of the firms' costs. The same comparison governs the slope of firm 2's best response function at the equilibrium associated with $t = 0$. In particular, this slope is $R_2'(q_1^*(0)) = (c_1 + c_2)/2c_2 - 1$. This is equal to zero when $c_1 = c_2$, in which case firm 1's degenerate conjecture happens to be consistent at the equilibrium, explaining why the optimal transfer price is exactly zero in this case. On the other hand, the slope is positive when $c_2 < c_1$, in which case the degenerate conjecture is inconsistently low and thus t^S is positive, while the opposite is true when $c_1 < c_2$.

[24] By contrast, if products are undifferentiated ($s = 1$), then $\nu^B = \widetilde{\nu} = -1$, and thus homogeneous Bertrand competition yields a CCE.

[25] See Dockner [1992] for a more detailed overview of the criticisms of conjectural variation models.

5 Distorted Games

In the previous section, we showed that the joint-profit maximizing level of double marginalization induces the same equilibrium behavior as a consistent conjecture, and that noncooperative double marginalization is therefore profitable when the upstream transfer price inadvertently induces the downstream firm to behave as if it has a more consistent conjecture. In this section, we introduce a generalized "distorted game" that nests a wide range of well-known behavior-distorting phenomena created by third parties, such as double marginalization, vertical restraints, or other mechanisms for strategic delegation. We synthesize each of these literatures with cooperative double marginalization by showing that every distortion always engenders downstream equilibria corresponding to consistent conjectures – Stackelberg equilibria or CCEs – even though all firms actually maintain ordinary degenerate Nash conjectures.

First consider a "generalized duopoly game" (GDG) based on the model from Section 3, but without double marginalization ($t = 0$). Each firm $i = 1, 2$ chooses a strategy $a_i \in A_i$ to maximize $\pi_i(a_i, a_j)$. Assume that assumptions (A1), (A2), and (A5) from Section 3 are satisfied for each i. In this game, profit maximization for firm i is characterized by the first order condition $\partial\pi_i(a_i, a_j)/\partial a_i = 0$, which pins down a continuously differentiable best response function $R_i(a_j)$. Now consider a *distorted game* based on the GDG. In this game, each firm's behavior is distorted (e.g. by double marginalization) via an alteration of its *internalized* payoff, but the firm still generates a total payoff of π_i. Formally:

Definition 1. *A* distorted game *is a perturbation of the generalized duopoly game such that each firm i's internalized payoff is altered in such a way that its first order conditions now take the form*

$$\frac{\partial \pi_i(a_i, a_j)}{\partial a_i} + \psi_i(a_i, a_j, \theta_i) = 0 \tag{17}$$

for some function ψ_i and some parameter $\theta_i \in \Theta_i \subset \mathbb{R}$.

We assume that ψ_i is continuously differentiable in (a_i, a_j, θ_i) and strictly monotonic in θ_i, and that Θ_i is compact. Here ψ_i is a "distortion function," which we interpret as specifying the *nature* of the distortion (e.g. double marginalization versus resale price maintenance). The scalar θ_i is a "distortion parameter" that captures the *magnitude* of the distortion (e.g. the specific price level mandated by a resale price maintenance contract). The best responses pinned down by (17) are assumed to be unique. As such, each i has a continuously differentiable best response function, denoted $\widetilde{R}_i(a_j|\theta_i)$, which generally differs from $R_i(a_j)$.

The distorted game nests many familiar situations in which downstream competition is influenced by third parties, a few of which are illustrated below. The general approach is similar to a model considered in Fabinger and Weyl [2013]. They consider the the class of symmetric quantity-choice games whose equilibria are characterized by conditions of the form $\epsilon_D L = \theta$, where L is the Lerner index, ϵ_D is demand elasticity, and θ is a "conduct parameter." In their model, different values of θ correspond to different games, such as conventional Cournot competition, a CV game, or Bertrand competition with differentiated products. Our distorted game is essentially a generalization of this approach, although we focus on a different interpretation of the comparative statics. Our interpretation is that the underlying game (the GDG) does not vary with θ, but rather this parameter reflects the extent to which behavior is being distorted. Consider a few well-known kinds of distorted games that fall within this general framework:

Competition with Conjectures

Although we are ultimately interested in Nash behavior, it will be useful to compare the actions of distorted Nash firms to CV firms, and thus we begin by showing that the distorted game nests a generalized CV game based on the GDG. Suppose each firm i maximizes π_i and has a (generally nonconstant) conjecture $\nu_i(a_i) = \partial a_j(a_i)/\partial a_i$ about firm j's response behavior.[26] In this case, profit maximization for firm i is characterized by

$$\frac{\partial \pi_i(a_i, a_j)}{\partial a_i} + \frac{\partial \pi_i(a_i, a_j)}{\partial a_j}\nu_i(a_i) = 0 \tag{18}$$

We will refer to an agent who behaves in this way as a "CV firm." A standard Nash firm will behave like the CV firm when the distortion function ψ_i takes the form

$$\psi_i(a_i, a_j, \theta_i) = \frac{\partial \pi_i(a_i, a_j)}{\partial a_j} f_i(a_i, \theta_i) \tag{19}$$

for some function f_i and parameter θ_i such that $f_i(a_i, \theta_i) = \nu_i(a_i)$.

Double Marginalization

Following Section 3, under double marginalization firm 1 must pay a total transfer price of the form $tX_1(a_1, a_2)$ to an upstream firm, where t is the transfer price parameter imposed by an upstream firm. To capture this in the distorted game, set $\theta_1 = t$ and

$$\psi_1(a_1, a_2, \theta_1) = -\theta_1 \frac{\partial X_1(a_1, a_2)}{\partial a_1} \tag{20}$$

Delegation to an Agent with Different Preferences

The choice of a_1 is delegated to an agent who maximizes some function $D(a_1, a_2, \theta_1)$, which differs from π_1. For example, similar to Sklivas [1987] and Vickers [1985], suppose that competition is in prices ($a_i = \text{price}_i$) and that the agent exaggerates production costs in its effort to maximize profits (and therefore sets systematically higher prices), so that D takes the form $D(a_1, a_2, \theta_1) = \pi_1(a_1, a_2) - \theta_1 C\big(q_1(a_1, a_2)\big)$, where C is a cost function, q_1 gives firm 1's output as a function of

[26] This is a generalization of the output-choice CV game in Bresnahan [1981].

(a_1, a_2), and $\theta_1 > 0$. Delegation of this sort is equivalent to a distorted game with distortion function

$$\psi_1(a_1, a_2, \theta_1) = -\theta_1 \frac{\partial C\big(q_1(a_1, a_2)\big)}{\partial q_1} \frac{\partial q_1(a_1, a_2)}{\partial a_1} \tag{21}$$

Resale Price Maintenance

Competition is in prices. A supplier employs a resale price maintenance strategy that requires firm 1 to set $a_1 = \theta_1$. Let

$$\psi_1(a_1, a_2, \theta_1) = -\frac{\partial \pi_1(\theta_1, a_2)}{\partial a_1} \tag{22}$$

for all (a_1, a_2). Then $\widetilde{R}_1(a_j|\theta_1) = \theta_1$ for every a_j, as required by the resale price maintenance contract.

Merger, Collusion, and Horizontal Shareholding

Suppose that each firm i chooses a_i to maximize $(1-\sigma_j)\pi_i(a_i, a_j)+\sigma_j\pi_j(a_j, a_i)$, where $\sigma_1, \sigma_2 \in [0, 1)$. If $\sigma_1 = \sigma_2 = \frac{1}{2}$, then the interpretation is that the firms have merged, or else that they are colluding. Otherwise the interpretation is that the firms are involved in horizontal shareholding, meaning that each firm i owns a fraction σ_i of total shares in its rival.[27] Let $\theta_i = \sigma_i/(1 - \sigma_j)$ for each i. Then this game's equilibrium arises in the distorted game when each ψ_i is defined by

$$\psi_i(a_i, a_j, \theta_i) = \theta_i \frac{\partial \pi_j(a_j, a_i)}{\partial a_i} \tag{23}$$

5.1 Strategic Distortions Induce Consistent Behavior

Now suppose the distorted game is the second stage in a sequential game. In stage 1, one or both of the distortion parameters θ_1 and θ_2 are chosen "strategically" by upstream firms. This means that a distinct upstream firm chooses θ_i to maximize $\pi_i\big(a_i^e(\boldsymbol{\theta}), a_j^e(\boldsymbol{\theta})\big)$, where $\big(a_1^e(\boldsymbol{\theta}), a_2^e(\boldsymbol{\theta})\big)$ denotes the downstream equilibrium generated by $\boldsymbol{\theta} \equiv (\theta_1, \theta_2)$, which is assumed to be unique and interior for every $\boldsymbol{\theta} \in \Theta_1 \times \Theta_2$. We can think of $\pi_i\big(a_i^e(\boldsymbol{\theta}), a_j^e(\boldsymbol{\theta})\big)$ as the equilibrium joint profits of firm i and the upstream firm (namely the one that chooses θ_i).

Since ψ_i is strictly monotonic in θ_i, it follows from (A2) that $\widetilde{R}_i(\cdot|\theta_i)$ shifts strictly monotonically in θ_i. Since $\widetilde{R}_j$ is independent of θ_i, this implies that $a_i^e(\boldsymbol{\theta})$ is strictly monotonic in θ_i. We extend assumption (A5) by assuming that, for every $\theta_j \in \Theta_j$, firm i's Stackelberg objective function $\pi_i(a_i, \widetilde{R}_j(a_i|\theta_j))$ is strictly concave in a_i.

Let $\widetilde{r}_i(a_j|\theta_i) \equiv \partial\widetilde{R}_i(a_j|\theta_j)/\partial a_j$ for each i. Recall that, in a CV game, firm i's conjecture about firm j is said to be consistent if it equals the slope of j's true best response function (this could be true globally or merely locally.) Thus, in the distorted game, $\widetilde{r}_j(\cdot|\theta_j)$ would be the (globally) consistent conjecture about firm j. With this, we can define a relationship between θ_i and θ_j that is analogous to the notion of conjectural consistency.

[27] Collusion via mutual shareholding has recently been investigated by Azar et al. [2015].

Definition 2. *θ_i is* consistent *with θ_j if*

$$\psi_i(a_i^e(\boldsymbol{\theta}), a_j^e(\boldsymbol{\theta}), \theta_i) = \frac{\pi_i(a_i^e(\boldsymbol{\theta}), a_j^e(\boldsymbol{\theta}))}{\partial a_j} \widetilde{r}_j(a_i^e(\boldsymbol{\theta})|\theta_j) \tag{24}$$

Following (18) and (19), this says that θ_i induces firm i to play the same equilibrium strategy it would adopt if it were a CV firm with the consistent conjecture $\nu_i(\cdot) = \widetilde{r}_j(\cdot|\theta_j)$. We therefore say θ_i induces "consistent behavior."

We now begin the equilibrium analysis by considering the case in which just one of the distortion parameters is chosen strategically in the first stage, while the other is left exogenous. (The possibility that one firm is not distorted at all is a special case.) Define ρ_i by

$$\rho_i(\theta_j) = \arg\max_{\theta_i \in \Theta_i} \pi_i\big(a_i^e(\boldsymbol{\theta}), a_j^e(\boldsymbol{\theta})\big) \tag{25}$$

Thus, $\rho_i(\theta_j)$ is the strategic choice of θ_i conditional on θ_j. Note that this maximizer is unique for every $\theta_j \in \Theta_j$.[28] We do not assume it is interior, however, because an upstream firm may be constrained in the extent to which it can distort a downstream firm. Our first result establishes that $\rho_i(\theta_j)$ is consistent with θ_j whenever the constraint on θ_i is non-binding.

Proposition 4. *If $\rho_1(\theta_2)$ is interior, then it is consistent with θ_2.*

Proof: Let $\boldsymbol{\theta} = (\rho_1(\theta_2), \theta_2)$. Since $\rho_1(\theta_2)$ is interior, it is pinned down by the first order condition

$$\frac{\partial \pi_1\big(a_1^e(\boldsymbol{\theta}), a_2^e(\boldsymbol{\theta})\big)}{\partial a_1}\left(\frac{\partial a_1^e(\boldsymbol{\theta})}{\partial \theta_1}\right) + \frac{\partial \pi_1\big(a_1^e(\boldsymbol{\theta}), a_2^e(\boldsymbol{\theta})\big)}{\partial a_2}\left(\frac{\partial a_2^e(\boldsymbol{\theta})}{\partial \theta_1}\right) = 0$$

$$\iff \quad \frac{\partial a_1^e(\boldsymbol{\theta})}{\partial \theta_1}\left[\frac{\partial \pi_1\big(a_1^e(\boldsymbol{\theta}), a_2^e(\boldsymbol{\theta})\big)}{\partial a_1} + \frac{\partial \pi_1\big(a_1^e(\boldsymbol{\theta}), a_2^e(\boldsymbol{\theta})\big)}{\partial a_2} \widetilde{r}_2(a_1^e(\boldsymbol{\theta})|\theta_2)\right] = 0$$

The bracketed sum must be equal to zero, since $a_1^e(\boldsymbol{\theta})$ is strictly monotonic in θ_1. It follows that

$$\psi_1\big(a_1^e(\boldsymbol{\theta}), a_2^e(\boldsymbol{\theta}), \rho_1(\theta_2)\big) = \frac{\partial \pi_1\big(a_1^e(\boldsymbol{\theta}), a_2^e(\boldsymbol{\theta})\big)}{\partial a_2} \widetilde{r}_2(a_1^e(\boldsymbol{\theta})|\theta_2) \tag{26}$$

Therefore $\rho_1(\theta_2)$ is consistent with θ_2. *Q.E.D.*

Thus a strategic distortion generates consistent behavior by the distorted firm. We hinted this result in Section 4 in the special case of double marginalization, but Proposition 4 shows that for any specification of ψ_1, the distortion parameter θ_1 will be chosen so that $\psi_1(\cdot, \theta_1)$ generates consistent behavior in equilibrium. This comports with the strategic delegation literature, which generally

[28] This follows from the facts that, for every $\theta_j \in \Theta_j$, $\pi_i(a_i, R_j(a_i|\theta_j))$ is strictly concave in a_i and $a_i^e(\boldsymbol{\theta})$ is strictly monotonic in θ_i.

suggests that different kinds of delegation practices are equivalent when used to strategically manipulate the behavior of a single downstream firm: they all induce the Stackelberg equilibrium. The following corollary provides a very general expression of this result.

Corollary 1. *If firm 2's behavior is not distorted ($\widetilde{R}_2(\cdot|\theta_2) = R_2$), then a strategic and interior choice of θ_1 generates the Stackelberg equilibrium of the GDG.*

This result is unsurprising in light of the strategic delegation literature. But the literature has not identified an analogous unifying principle for equilibria in which *both* downstream firms are strategically distorted by upstream competitors. This is perhaps due to a conceptual difficulty: a single distortion makes the distorted firm act like a Stackelberg leader, but what does it mean for both downstream firms to act like Stackelberg leaders? We submit that conjectural consistency – and in particular *mutual* consistency – provides the right theoretical machinery for thinking about this. To that end, now suppose the two-stage game involves *competing distortions*: θ_1 and θ_2 are both chosen strategically (and simultaneously) by upstream competitors in stage 1. That is, for each i, a distinct upstream firm chooses θ_i to maximize $\pi_i\big(a_i^e(\boldsymbol{\theta}), a_j^e(\boldsymbol{\theta})\big)$, taking θ_j as given. We refer to an equilibrium of this game as a *competing distortions equilibrium*, denoted $\big(\boldsymbol{\theta}^*, a_1^e(\boldsymbol{\theta}^*), a_2^e(\boldsymbol{\theta}^*)\big)$. This game is a generalization of many well-known models from the strategic delegation literature, such as Bonanno and Vickers [1988], Rey and Stiglitz [1995], Jansen et al. [2007], Fershtman and Judd [1987] and Gal-Or [1991], among others. Consider the following definition:

Definition 3. *The profile (θ_1, θ_2) generates an* induced consistent conjectures equilibrium *(induced CCE) if θ_1 and θ_2 are mutually consistent.*

That is, an induced CCE is an equilibrium in which each downstream firm is induced to play the same equilibrium strategy it would adopt if it had a (globally) consistent conjecture about its rival's (distorted) response behavior. This is a natural analogue to the CCE concept from Bresnahan (1981) although, as discussed below, these concepts do not always produce identical equilibrium strategies. Our next result is that competing distortions generate an induced CCE whenever they are interior, which is an immediate implication of Proposition 3.

Proposition 5. *Let $\big(\boldsymbol{\theta}^*, a_1^e(\boldsymbol{\theta}^*), a_2^e(\boldsymbol{\theta}^*)\big)$ be a competing distortions equilibrium. If $\boldsymbol{\theta}^*$ is interior, then it generates an induced CCE.*

This link between Nash equilibria and the CCE concept is surprising. A common criticism of conjectural variation analysis is that, when conjectures are supplied artificially by the modeler (the CCE being a special case), the resulting equilibria are in general not Nash equilibria of the underlying game. In the CCE case, the firms' behavior appears to correspond to a dynamic interaction that is not actually being modeled. Accordingly, CCEs do not ordinarily arise organically. But Proposition 5 shows that this is not so when firms are strategically distorted. On the contrary, in this case CCEs naturally arise in Nash equilibria. In this situation, Bresnahan's non-Nash CCE concept can serve as a reduced form analysis of a multi-stage interaction in which Nash firms are strategically distorted in parallel. This also helps to provide a broader intuition for what is happening in the strategic delegation literature.

Our induced CCE and the CCE concept defined in Bresnahan [1981] – henceforth referred to as the *Bresnahan CCE* – do not always produce exactly the same downstream equilibrium strategies, although competing strategic distortions will generally move downstream strategies in the direction of the Bresnahan CCE.[29] In an induced CCE, each firm behaves *in equilibrium* as if it has a consistent conjecture, but its beliefs may be inconsistent elsewhere, i.e. at other points along its rival's reaction function. This is reflected in our consistency definition, which hinges only on the value taken by ψ_i in equilibrium. But, as briefly noted in Section 4, the Bresnahan CCE requires that each firm's conjecture is consistent *in an open neighborhood* of the equilibrium. (Bresnahan does not explicitly state his reasoning for imposing this stronger requirement.[30]) A given distortion function ψ_i may lack the flexibility to achieve consistency throughout an open neighborhood of the equilibrium, which may prevent the game from producing the same downstream strategies as a Bresnahan CCE. We illustrate this point in Examples 6A and 6B, below.

Nonetheless, the induced CCE exhibits the principal properties that distinguish the Bresnahan CCE. First, both concepts always have the property that, for each i, the equilibrium is the π_i-maximizing profile along j's reaction curve.[31] Second, the first property does not arise merely by coincidence. Assuming that each firm's payoff function is strictly monotonic in its rival's strategy (at least at the equilibrium point), then in an ordinary (undistorted) Nash game, one of the following statements will always be false for each i: (1) the equilibrium is i's favorite profile along j's reaction function; or (2) firm j's reaction function has nonzero slope at the equilibrium. Property (1) can be true, as in the iso-elastic Cournot game (Example 2) with $c_1 = c_2$, but only if i's degenerate Nash conjecture simply happens to be consistent at the equilibrium, i.e. only if the rival's reaction function has slope zero at the equilibrium. But this is pure coincidence; it does not reflect strategic foresight. Alternatively, if (2) is true, then i can benefit by moving slightly up or down j's reaction curve, and so (1) must be false. However, in both an induced CCE and a Bresnahan CCE, both (1) and (2) are generally true,[32] reflecting equilibrium play that is mutually-consistent due to strategic anticipation, not by coincidence.

To reconcile the induced CCE with Bresnahan (1981), we define a stronger notion of consistency between the distortion parameters θ_1 and θ_2. We say that θ_i is *perfectly consistent* with θ_j if there exists $\varepsilon > 0$ such that ψ_i satisfies

$$\psi_i(a_i, \widetilde{R}_j(a_i|\theta_j), \theta_i) = \frac{\partial \pi_i\big(a_i, \widetilde{R}_j(a_i|\theta_j)\big)}{\partial a_j} \widetilde{r}_j(a_i|\theta_j) \tag{27}$$

[29] For example, in a Cournot game, the Bresnahan CCE will involve higher output levels than the Nash equilibrium, and so will any induced CCE.

[30] On page 936, Bresnahan [1981] notes that this ensures that each firm i has correct beliefs about all higher-order derivatives of $\widetilde{R}_j$ at the equilibrium point – not just the first derivative – although he does not clarify why he deems this essential.

[31] In Bresnahan [1981], this need only be true in an open neighborhood along the rival reaction curve.

[32] The only exception is that (2) can in principle still be false, in which case the strategic distortion happens to involve no distortion at all, since the Nash conjecture happens to be correct.

for all a_i such that $|a_i - a_i^e(\boldsymbol{\theta})| < \varepsilon$. This says that firm i is acting like it has correct beliefs about $\widetilde{r}_j(\cdot|\theta_j)$ at all a_i-values within an open neighborhood of its own equilibrium strategy. This is a direct analogue to the consistency definition in Bresnahan [1981]. We then say that $\boldsymbol{\theta}$ generates an *induced perfect CCE* if θ_1 and θ_2 are mutually perfectly consistent. An induced perfect CCE therefore yields a downstream equilibrium corresponding to a Bresnahan CCE. However, competing distortions may or may not generate an induced perfect CCE, because some ψ_i functions can achieve consistency only at one point, which will always be the equilibrium. We illustrate this with two short examples, only one of which generates an induced perfect CCE. Both examples are based on the linear Cournot game with differentiated products that underpinned the linear CV game from Section 4.

Example 6A (Competing Linear Transfers): Consider the linear quantity-choice CV game introduced in Section 4, but now assume the firms have ordinary degenerate conjectures. Recall that payoffs are $\pi_i(q_i, q_j) = (1 - q_i - sq_j)q_i$ for each i, where $s \in (0, 1)$. Now suppose each i is double marginalized according to the linear transfer price schedule $\theta_i X_i(q_i, q_j) = \theta_i q_i$, so that $\psi_i(q_i, q_j, \theta_i) = -\theta_i$. This yields a reaction slope $\widetilde{r}_i(q_j|\theta_i) = -s/2$, which is independent of θ_i. Therefore each upstream firm will strategically induce its downstream partner to behave as if it has a conjecture of $-s/2$, which is indeed consistent with its rival's behavior. However, this induced CCE is still not perfect. We can see this indirectly by noting that $-s/2 \neq \nu^\star$, where $\nu^\star = -(1 - \sqrt{1 - s^2})/s$ is the Bresnahan CCE conjecture from the linear CV game based on this Cournot framework (see Section 4). To see it directly, it can be verified that the perfect consistency condition (27) is $\theta_i = -\frac{1}{2}s^2 q_i$ in this model, and clearly there is no θ_i-value such that this will be satisfied throughout an open neighborhood of q_i-values. Thus it will be satisfied only in the equilibrium. ■

Example 6B (Competing Quadratic Transfers): Amend the last example so that each transfer price schedule takes the quadratic form $\theta_i X_i(q_i, q_j) = \frac{1}{2}s\theta_i q_i^2$, and thus $\psi_i(q_i, q_j, \theta_i) = -s\theta_i q_i$ for each i. This yields $\widetilde{r}_i(q_j|\theta_i) = -s/(2 + s\theta_i)$, so the upstream firms can affect reaction function slopes in this case. The perfect consistency condition (27) reduces to

$$\theta_i = \frac{-s}{2 + s\theta_j} = \widetilde{r}_j(\cdot|\theta_j) \tag{28}$$

Since this equation is invariant in q_i, this says that θ_i is either always consistent with θ_j, or else never consistent. Of course, we know it is consistent at the equilibrium level of q_i, and thus it must be consistent everywhere. Hence the induced CCE is perfect. We could see this indirectly by noting that symmetry implies both firms will choose θ^* such that $\theta^* = \widetilde{r}_j(\cdot|\theta^*)$, and solving this yields $\theta^* = \nu^\star$. This tells us that each firm will be induced to act as though it maintains the Bresnahan CCE conjecture $\nu^\star$. ■

The intuition is the following: as already noted, the downstream behavior underpinning the

Bresnahan CCE can be replicated in the distorted game through the appropriate choice of distortion function. Similarly, each possible distortion mechanism (e.g. resale price maintenance) corresponds to a particular kind of distortion function. Some of these are equivalent to the CV distortion function – the implicit specification of ψ_i in a CV game – as in Example 6B, while others are distinct, as in Example 6A. By "distinct," we mean that it distorts the downstream reaction function $\widetilde{R}_i(a_j|\theta_i)$ in a different way – as θ_i increases, it elicits different measures of "shifting" and "twisting." Thus, a change in ψ_i may in turn change the strategic choice of θ_j, since upstream firm j is now choosing its favorite profile along a different reaction function. The result is that the change in ψ_i ultimately changes the downstream equilibrium. However, competing distortions generally move downstream strategies in the direction of the Bresnahan CCE profile[33].

6 Discussion and Concluding Remarks

Traditional analysis of double marginalization with successive monopolies suggests that the firms have an incentive to eliminate it.[34] We have shown that this intuition does not hold when downstream firms face competition, whether or not the vertically related firms act cooperatively to maximize joint profits. In particular, we provide four novel results which fully describe how double marginalization operates with downstream oligopoly, both when upstream and downstream firms coordinate and when they do not.

First, even if the firms interact only through noncooperative linear pricing, and even if the downstream rivals' costs are left exogenous, double marginalization may inadvertently increase joint profits in many familiar economic scenarios. Second, when double marginalization enhances profits, it does so by creating a beneficial strategic effect in rival behavior that outweighs the negative direct effect of distorting the downstream firm's behavior, essentially approximating strategic delegation à la Bonanno and Vickers [1988]. The relationship between these effects is ambiguous and amenable to analysis using traditional price theoretic tools. Third, many behavior-distorting phenomena created by third party decision-making, such as double marginalization, vertical restraints, and strategic delegation are intimately related. We unify such practices in a generalized "distorted game" which demonstrates that jointly-optimal (i.e. strategic) distortions, no matter their exact nature, are always those which induce downstream firms to behave as if they have consistent conjectures about rival behavior. Fourth, if downstream firms are strategically distorted in parallel by upstream competitors, then these distortions generate an "induced consistent conjectures equilibrium," a concept that closely analogizes – and sometimes coincides with – the CCE concept from Bresnahan [1981]. This is surprising since the induced CCE arises from Nash behavior, whereas Bresnahan's CCE concept is generally regarded as non-Nash. It suggests that the CCE concept can be interpreted as the reduced form of a multi-stage Nash equilibrium involving competitors who are being strategically distorted.

[33]This is clearly true if the distorted game is supermodular, for in this case the nature of a strategic distortion (whether a_i should be increased or decreased) will be the same for all ψ_i.

[34]Riordan and Salop [1995] provide an excellent discussion of antitrust perspectives on double marginalization.

That double marginalization can enhance joint profits when the downstream market involves a strategic interaction embodies a more general principle that is already known: an agency problem can be beneficial to the extent that it facilitates beneficial commitment and thereby induces a more desirable equilibrium. This is the logic that underpins the theory of strategic delegation. As Vickers [1985] pointed out in his seminal paper on delegation, "[i]f control of my decisions is in the hands of an agent whose preferences are different from my own, I may nevertheless prefer the results to those that would come about if I took my own decisions. This has some interesting implications for the theory of the firm. For example, in markets where firms are interdependent, it is not necessarily true that maximum profits are earned by firms whose objective is profit-maximisation." However, prior studies applying the idea of strategic delegation to vertical relationships differ from our noncooperative analysis in an important respect, namely that they require delegation to be a credible form of commitment.[35] The possibility of credibly committing to non-Nash actions by using two-part tariffs, manager delegation, or similar techniques has been widely criticized (e.g., Katz [1991], Corts and Neher [2003], O'Brien and Shaffer [1992]).

Our results on noncooperative double marginalization are not subject to this credibility problem. The firms in our model are, by assumption, unable to coordinate. They cannot commit to contracts designed to maximize joint profits by taking advantage of strategic effects, as reflected by the upstream firm's inability to enter a non-renegotiable two-part tariff contract. Nor is there any one agent that has the power to unilaterally "overrule" the noncooperative price levels at the last moment. Thus, in our noncooperative framework, the firms' misaligned incentives are not the product of any strategic decisions to partition or manipulate control, but rather are an incidental consequence of the firms' inability to coordinate, which is a constraint generated by exogenous phenomena – e.g. transaction costs, hold-up, or other sources of market failure. The flipside of this exogenous misalignment is that the profit effects of double marginalization are imperfect – it generally falls short of maximizing joint profits, and in many environments it erodes them. Put simply, noncooperative double marginalization operates like more credible yet less precise strategic delegation.

The size and scope of the firm, then, can depend on indirect strategic effects even if we accept the "credibility critique" that firms who can coordinate cannot commit to behaving differently than an integrated firm. If the firms are unable to coordinate efficiently when separate (i.e. to write a contract that will maximize joint profits), they may nonetheless prefer not to integrate when noncooperative double marginalization induces the downstream firm to act less competitively, increasing joint profits via an indirect strategic effect. Unlike production externalities which can

[35] Among many such studies, Fershtman and Judd [1987] and Sklivas [1987] discuss who exactly should be delegated to, with the latter in particular pointing out that Cournot firms should delegate to price-setters who place extra weight on revenues (thereby understating costs) in their effort to maximize profits. Alles and Datar [1998] gives a similar application in an operations/strategic management context, and Jansen et al. [2007], extends the idea to delegates maximizing market share rather than revenue alone. Rey and Stiglitz [1995] consider a model in which upstream competitors compete "through" downstream agents (their distributors), using two-part tariffs to strategically distort downstream pricing by manipulating the agent's costs. For formal investigations of what results principals can achieve by playing games "through" delegated agents, see Katz [1991] and Rustichini and Prat [2003].

be mitigated either through Coasean bargaining or integration, the "positive externality" created incidentally by an upstream markup exists only as a result of noncooperation; it would be eliminated by vertical integration or by the ability to bargain efficiently. Therefore, the critique that strategic delegation is noncredible does not imply that vertically related firms will always integrate to avoid being double marginalized.

References

Michael Alles and Srikant Datar. Strategic transfer pricing. *Management Science*, 44(4), 1998.

Jose Azar, Martin C. Schmalz, and Isabel Tecu. Anti-competitive effects of common ownership. *Ross School of Business Paper No. 1235*, 2015.

Michael Baye, Keith J Crocker, and Jiandog Ju. Divisionalization, franchising, and divestiture incentives in oligopoly. *American Economic Review*, 86(1), 1986.

GIacomo Bonanno and John Vickers. Vertical separation. *Journal of Industrial Economics*, 36(3), 1988.

A.L. Bowley. *The Mathematical Groundwork of Economics*. Oxford University Press, 1924.

Timothy F. Bresnahan. Duopoly models with consistent conjectures. *American Economic Review*, 71(5), 1981.

Kenneth S. Corts and Darwin V. Neher. Credible delegation. *European Economic Review*, 47(3), 2003.

Engelbert J. Dockner. A dynamic theory of conjectural variations. *The Journal of Industrial Economics*, 1992.

Michal Fabinger and E. Glen Weyl. Pass-through as an economic tool: Principle of incidence under imperfect competition. *Journal of Political Economy*, 121(3):528–583, 2013.

Chaim Fershtman and Kenneth Judd. Equilibrium incentives in oligopoly. *American Economic Review*, 77(5), 1987.

Esther Gal-Or. Duopolistic vertical restraints. *European Economic Review*, 1991.

Géarard Gaudet and Ngo Long. Vertical integration, foreclosure, and profits in the presence of double marginalization. *Journal of Economics & Management Strategy*, 5(3):409–432, 1996.

Marhijs Jansen, Arie van Lier, and Arjen van Witteloostuijn. A note on strategic delegation: The market share case. *International Journal of Industrial Organization*, 25(3), 2007.

Michael L. Katz. Game-playing agents: Unobservable contracts as precommitments. *RAND Journal of Economics*, 22(3), 1991.

Frank Mathewson and Ralph Winter. An economic theory of vertical restraints. *RAND Journal of Economics*, 1984.

Daniel P. O'Brien and Greg Shaffer. Vertical control with bilateral contracts. *RAND Journal of Economics*, 23(3), 1992.

Tonu Puu. Chaos in duopoly pricing. *Chaos, Solitons and Fractals*, 1(6):573–581, 1991.

Patrick Rey and Joseph Stiglitz. The role of exclusive territories in producers' competition. *RAND Jour*, 26(3), 1995.

Michael H. Riordan. Vertical integration. In Steven N. Durlauf and Lawrence E. Blume, editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, 2nd edition, 2008.

Michael H. Riordan and Steven C. Salop. Evaluating vertical mergers: A post-chicago approach. *Antitrust Law Journal*, 63(2):513–568, 1995.

Aldo Rustichini and Andrea Prat. Games played through agents. *Econometrica*, 71(4), 2003.

Steven C. Salop and David T. Scheffman. Raising rivals' costs. *The American Economic Review Papers & Proceedings*, 73(2), 1983.

Steven D. Sklivas. The strategic choice of managerial incentives. *RAND Journal of Economics*, 18 (3), 1987.

Joseph J. Spengler. Vertical integration and antitrust policy. *Journal of Political Economy*, 58(4), 1950.

John Vickers. Delegation and the theory of the firm. *The Economic Journal*, 95, 1985.

Oliver E. Williamson. The vertical integration of production: Market failure considerations. *The American Economic Review Papers & Proceedings*, 61(2), 1971.