Essays in Industrial Organization

by

Jenny Rae Hawkins

A Dissertation Submitted to the Faculty of the

Department of Economics

In Partial Fulfillment of the Requirements
For the Degree of

Doctor of Philosophy

In the Graduate College

The University of Arizona

2 0 1 1

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Jenny Rae Hawkins
entitled Essays in Industrial Organization
and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy

____ Date: 05/16/2011
Jonah B. Gelbach

____ Date: 05/16/2011
Stanley S. Reynolds

____ Date: 05/16/2011
Katherine Y. Barnes

____ Date: 05/16/2011
Price V. Fishback

____ Date: 05/16/2011
Keisuke Hirano

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.
I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

____ Date: 05/16/2011
Dissertation Director : Jonah B. Gelbach

____ Date: 05/16/2011
Dissertation Director : Stanley S. Reynolds

Statement by The Author

Signed: Jenny R. Hawkins

Acknowledgments

I extend my gratitude to many genuine people who supported me throughout my graduate studies. Dr. Jonah Gelbach, my advisor, offered endless support and encouragement at all stages of my research. As his assistant and student, many fruitful, if not entertaining discussions about economics broadened my knowledge, interests, and pursuits, and I am deeply grateful. I appreciate my other advisor, Dr. Stan Reynolds, for many discussions of my work, valuable insight about future research, and advice whenever requested. Dr. Kei Hirano went beyond his call of duty as a committee member, and I am grateful for his gentle guidance and extra time. Since my first year, Dr. Price Fishback's feedback and continuous encouragement have helped boost my confidence and excitement for research. Dr. Kathie Barnes provided fresh perspectives on my research, as well as appreciated advice as a mentor.

Dr. Gerry Swanson's passion for economics is contagious. I thank him and Dr. Don Wells for allowing me to learn year after year from them. Dr. John Drabicki's compassion, balanced with just enough sternness as a teacher, are exemplary, and I am thankful to have always received from the former side of the scale. I value Dr. Tracy Regan's friendship and advice. Lana, you supported me beyond what you promised; thank you. To my undergraduate advisor, Dr. Cynthia Hill, who opened my eyes to economics, and other faculty in our department who offered guidance, thank you.

I could not have asked for a better cohort and set of friends, especially during the last few years of graduate school. Natalia has been a true friend who constantly encourages me. I enjoyed many useful research discussions with Matt, who kept me amused as an office mate and friend. Trevor's constant interruptions will be missed and in fact made the long work days less mundane; he was a great office mate. With half of our cohort being left-handed, I must acknowledge my fellow lefties and true friends, Eylem, Jon, Trevor, and Steven. I also appreciate the sincere friendships of Dina, Theresa, Keoka, Amilcar, Dong-Hyuk, and Taylor.

The sacrifices and encouragement of my family helped me endure. I thank my mom, Joan, Shelli, my immediate family and Gordon, Tom, Jean and Annabelle, as well as all of the other Manos' who treat me as their own. My mom, Jean and Shelli constantly lifted my spirits and ordered me to take a break when I failed to see the value in one. Everyone offered support and love through Tucson visits and comforting phone calls. I greatly appreciate the Pickrells, who graciously shared their part of this beautiful desert with me. Sunsets near the Catalinas and hiking among the saguaros never were often enough, even if they resulted in a bothersome bout of valley fever. My time here included life's highest of highs and lowest of lows. I will cherish many memories and the opportunity to study economics at The University of Arizona.

Dedication

> "It's the ultimate test of the human spirit to see who can find the most in them to keep on going, even when you're hurting. You don't want to finish the race and feel like, 'Man, I could have tried a little bit harder or ran a little bit faster.' You want to finish the race and feel like, 'I gave it all I had.'" *–Steven Michael Manos, 1981-2008*

In memory of

> my dad, Dale; as a little girl I solved math problems to be like you, and this is where it led me,

and of course,

> Steven. You always will hold my heart. The endless hours at our cubicles and the library studying and exploring economics together was worth every minute, if only to share the time I had with you. This endeavor has been as much for you as for myself.

Table of Contents

List of Figures

List of Tables

ABSTRACT

This dissertation consists of three essays evaluating topics in industrial organization. The first essay investigates a market structure or property regime in which a final good exists only by assembling multiple, monopoly-supplied components. In such dynamic settings, any sunk cost results in an outcome of hold-up, also known as tragedy of the anticommons. I design a model showing conditions for which two factors that reduce sunk cost, refunds and complementarities, mitigate hold-up. If the first component purchased has positive stand alone value or the first seller offers a full refund, hold-up is mitigated. My results suggest several policies that can mitigate inefficient outcomes in assembly problems, including legal requirements on full refunds, regulation on the purchasing order of components, and prohibition of price discrimination. The second essay applies Bayesian statistics to single-firm event studies used in securities litigation and antitrust investigations. Inference based on Bayesian analysis does not require an assumption of normality that potentially invalidates standard inference of classical single-firm event studies. I investigate ten events, five from actual securities litigation cases. Various Bayesian models, including replication of the frequentist approach, are examined. A flexible Bayesian model, replacing parametric likelihood functions with the empirical distribution function, also is explored. Our approach suggests an alternative, valid method for inference with easy implementation and interpretation. The third essay, motivated in the context of pharmaceutical advertising, analyzes demand rotations caused by an exogenously determined advertising parameter under Cournot oligopoly competition. We find that firms and consumers prefer extreme levels of advertising, but preferences for which extreme do not necessarily align. However, these differences can be alleviated with few or many firms in the market or cheap or expensive technologies. Therefore, advertising levels, regulated or not, might not serve consumers' best interests unless certain market attributes hold.

Chapter 1

Mitigating Hold-up Through Complementarities and Refundability

1. Introduction

Consider a market structure where the assembly of multiple component goods creates a welfare-improving composite good. The components are complementary: only the combination of these multiple components creates the composite. Also, no substitute exists for any component. Each component is monopoly-supplied, and the monopolists do not collude. I refer to this market for the composite good as a "fragmented market" in which the potential market failure creates the "assembly problem."

Generally, the component and composite goods are bundles of rights, or properties, such as any physical good, creative work, or intangible right to use. Examples of a fragmented market setting include a developer trying to assemble the real property of multiple, distinct land owners; shared owners of a single piece of property (such as heirs), each with veto rights to exclude; or an innovator whose success depends on acquiring multiple licenses from unique patent holders.

The most common terminology used to analyze these fragmented markets is *complementary monopoly* (e.g. Sonnenschein (1968), Machlup and Taber (1960)), *complementary oligopoly* (e.g. Salinger (1989), Parisi and Depoorter (2003), and Dari-Mattiacci and Parisi (2006)), *anticommons property* (e.g. Heller (1998) and Buchanan and Yoon (2000)), and *gridlock* (Heller (2008)). Although the terminology varies by field of study and types of rights considered, the potential market failure is well-known. Because each component is controlled by separate monopolists, each supplier does not internalize the total benefit from the composite good. The fragmentation of rights can result in a total equilibrium price for the composite good higher than even that chosen by a vertically integrated monopolist, yielding suboptimal or no acquisition of the composite good. This inefficient outcome has been referred to as the tragedy of the anticommons, double marginalization, or hold-up.

I seek to understand the underlying mechanism driving the inefficient outcome and how to alter that mechanism so hold-up might be mitigated. In particular, I model this assembly problem in terms of sunk cost. Recognizing how sunk cost drives the inefficient outcome, I then analyze two factors that reduce sunk cost as solutions to overcome hold-up. The first factor, imperfect complementarities of component goods, corresponds to a component's residual benefit from the production technology. Second, refundability, corresponds to an alteration in the contract structure through the return and refund of a previously purchased component. These two factors reveal aspects of the problem not fully characterized in the literature, which itself has not explicitly recognized the role of sunk costs. Using the subgame perfection solution concept, I consider twelve combinations of imperfect complementarities and refundability in a sequential game of two sellers and one representative buyer.

My analysis generalizes a common assumption in models of fragmented markets that component goods are perfectly complementary. In this setting, perfect complementarities means no residual benefit exists in the production technology of any component. Therefore, if components are perfectly complementary, no component has value except in combination with the other required components. If no single component has value in use alone, then once purchased, the full cost of that component is sunk. Knowing the cost of every previously-purchased component is sunk, the last seller rationally responds by pricing without regard to the other components' prices. This results in a total equilibrium price so high that an inefficient quantity of composites – even fewer than sold by a single monopolist – is purchased in equilibrium.

However, monopoly-supplied components are not necessarily perfectly complementary. A land developer desiring to assemble multiple parcels of property might value one of the required parcels even if he is unable to assemble the property via acquisition of all other parcels. Therefore, a natural next question is whether *imperfectly* complementary components overcome the inefficient effects of sunk cost. An imperfectly complementary component has *stand alone value*, or value to the buyer other than in use to create the composite good. Because a component with stand alone value has outside use to the buyer, then the sunk cost to purchasing that component is reduced by that stand alone value. When complementarities are imperfect, the last seller must account for the possibility that potential buyers can use previously purchased components, even without the last seller's good. This fact puts downward pressure on the last seller's optimal price.

In my model, the composite good requires two components. I proceed by varying the degree of complementarity of each of these two components. If at least the first component supplied is imperfectly complementary to the second component supplied, no anticommons tragedy results. Consider the land developer seeking to assemble two distinctly-owned parcels of land, at least one of which can be resold or has use besides creating the composite parcel. As long as the developer purchases a parcel with positive stand alone value first, the last seller has no ability to hold-up the developer. Note that components are purchased in either a predefined or undefined order. This result also suggests that hold-up in fragmented markets can be induced if a regulation requires a particular order of assembling rights. If the order of acquiring rights optimally is chosen by the buyer or optimally is determined by a social planner, hold-up can be mitigated.

On the other hand, the last component purchased might be the only component with stand alone value. In this case, sunk cost remains full, as in the case of perfect complementarities. The last seller still holds up the buyer, and therefore, stand alone value in the last component does nothing to mitigate the anticommons tragedy. This intuitive result is unambiguous in a model where the capacity of the last component is constrained to the exact proportion required to create the composite good. For example, in land or permit acquisitions, capacity for each component is restricted to one unit, where the buyer chooses quantity one (buy) or quantity zero (not buy).

However, the last seller might be capable of supplying units of the component beyond that required to create the composite good. I show hold-up no longer necessarily results if (i) the last seller can supply a quantity beyond that required to create the composite and (ii) the buyer values that excess quantity in use alone. In this case, the degree of stand alone value, precisely defined below, determines whether hold-up results. Pricing incentives change when the last seller supplies not only the required number of units to create the final good, but also additional units for stand alone use. First, if the last component has low stand alone value, full hold-up results. Moreover, if the last seller can price discriminate between the good used alone and the good used to assemble the composite, hold-up results, regardless of the last component's degree of stand alone value. However, if the last component has high stand alone value and sellers cannot price discriminate, the anticommons tragedy is mitigated.

These outcomes result when the buyer must purchase the components in a particular order: the first component has no stand alone value and the last component has positive stand alone value. Therefore, even when a natural order to acquiring rights exists and cannot be changed, and even if the last component has stand alone value, hold-up does not necessarily result; without capacity constraints, high stand alone value in the last component mitigates hold-up, provided the last seller does not price discriminate.

Another factor that reduces sunk cost is a refund for previously purchased components. With a partial or full refund, the buyer has the added option of returning the first component. To my knowledge, no previous research considers refundability. If the first component has imperfect complementarities, the sunk cost from purchasing that component is reduced by that component's stand alone value. Refunds, on the other hand, reduce sunk cost by the amount of the recovered purchase price. While both factors reduce sunk cost, they alter the buyer's payoffs differently. With imperfect complementarities in the first component, the sunk cost can be fully eliminated, depending on the first seller's best response to the last seller's optimal behavior. With refunds, sunk cost is fully eliminated only if the refund is full, regardless of the first seller's optimal price.

From the last seller's perspective, when a positive refund is offered, optimal behavior might require not only ensuring the buyer purchases the last component, but also ensuring the buyer does not return the first component. However, the last seller knows if the first component is returned and the refund is not one hundred percent, the buyer's surplus is negative, because sunk cost remains. Therefore, if the refund is only partial or zero, then the reduction in sunk costs is not enough to give the last seller incentive to account for the effects of a refund. With a zero or partial refund, unless the first component also has stand alone value, full hold-up results with no purchase of the composite good in equilibrium. However, a full refund completely eliminates sunk cost, forcing the last seller to price low enough so the buyer purchases the last component and does not return the first component. Therefore, a full refund ensures that hold-up by the last seller is not privately optimal, guaranteeing a Pareto improving outcome.

My analysis also includes joining the factors of imperfect complementarities and refundability. When the first component has zero stand alone value, which corresponds to the case of perfect complementarities or the case of imperfect complementarities in the last component, I show only a full refund mitigates hold-up (assuming transaction costs are low enough). Recall, stand alone value in the first component mitigates hold-up. Therefore, allowing for refundability does not alter the outcome. However, I provide a welfare analysis, endogenizing the refund level, as well as the order of play, to show how equilibrium payoffs are affected by the level of refund. I show the first seller does not necessarily prefer a full refund. Also, depending on the degree of stand alone value, a first-mover or second-mover advantage exists, suggesting the last seller in fragmented markets does not always receive the highest equilibrium payoff.

My model reveals the role of sunk cost in assembly problems created by anticommons property or complementary monopoly. The results imply situations where complementarities or refundability can and cannot prevent hold-up. One policy response to market power in such settings might include divestiture of an anticompetitive (or potentially anticompetitive) firm. However, fragmenting a market by forcing sell-offs or spin-offs of a vertically-integrated firm into multiple entities might result in a total composite price exceeding the original firm's price, even when it is already the anticompetitive monopoly price. Therefore, the resulting tragedy of the anticommons could yield a decrease in social welfare. Policy prescriptions accounting for complementarities and refundability in such situations might avert a further decrease in efficiency. My results suggest policies, such as legal requirements on full refunds, regulation on order in which components must be purchased, or prohibition of price discrimination, to mitigate inefficient outcomes in assembly problems.

1.1. Related Literature

Many have considered markets for a composite good that cannot exist without assembling multiple, monopoly-supplied components. My contribution rests in motivating the assembly problem through sunk cost, then considering how complementarities and refundability affect sunk cost in a way that might mitigate hold-up. While some bargaining models of hold-up, such as Carmichael and MacLeod (2003) and Gul (2001), include discussion of sunk cost, to my knowledge, no models of the assembly problem explicitly model sunk costs.[1]

The seminal analyses of assembly problems in fragmented markets are due to Cournot (1838) and Ellet (1839). Cournot's examination, most often referred to as complementary monopoly, initially assumes perfectly complementary, monopoly-

[1] Llanes and Trento (2009) incorporate a fixed (sunk) cost in the production of components in their model of optimal patent policy, while in my model, the consumer's sunk cost drives the assembly problem.

supplied components with zero costs to produce the components and composite. Assuming simultaneous interaction between monopolists, Cournot concludes that the equilibrium total price is greater than the price charged by a single monopolist. He also shows that as the number of components required to assemble the composite good increases, so does the difference between this equilibrium total price and the vertically integrated price. Sonnenschein (1968) takes insight from Edgeworth (1925) to show how Cournot's theory of complementary monopoly is the dual of his more well known theory of duopoly.

While the complementary monopoly model in Cournot (1838) chooses prices simultaneously, Spengler (1950) considers a similar problem in a sequential setting, *a la* Stackelberg (1934). In both frameworks, the price for the composite increases with the number of required components. In fact, Feinberg and Kamien (2001) show in a game of perfect and complete information an outcome of hold-up is the analog result to double marginalization under a game of imperfect information. Thus, while hold-up results under sequential play, double marginalization results under simultaneous play. In both, the tragedy of the anticommons is the same; the quantity (price) of the composite good purchased in equilibrium is lower (higher) than under vertical integration.

The second seminal work on the assembly problem, Ellet (1839), reaches the same conclusion as Cournot (1838) with respect to total price and welfare. However, Ellet's motivation is that of trade (and tolls, related thereto). Acquiring permits to assemble a composite good motivates many analyses of this market structure. Feinberg and Kamien (2001) and Gardner, Gaston, and Masson (August 2002) consider the assembly of permits purchased from successive monopolists in the form of a toll. In this case, the composite good is a destination or privilege. This perspective provides a clear transition from consumption goods to those providing a "right to use" (such as a permit to traverse real property).

Ellickson (1993) and Fennell (2010) credit the inception of the term *anticommons* to Frank Michelman. Michelman (1982) describes the problem as the converse of a commons problem, calling it a property regime in which everyone has the right to block use and no one has the exclusive privilege of use.[2]

Heller (1998) brought this assembly problem to greater attention among legal scholars. Buchanan and Yoon (2000) follow up on Heller's account to show the symmetry between the tragedies of the commons and anticommons. Though a less general model than Cournot's, they also show total equilibrium price increases with the number of required components.

Refundability plays an important role in my analysis as the second factor I consider to reduce sunk cost and possibly mitigate hold-up. Buchanan and Yoon (2000) state a full refund is allowed; however, refunds play no role in their analysis. To my knowledge, no research considers refunds in the assembly problem.

Most theoretical and experimental models of complementary monopoly and anticommons property assume perfect complementarities among the component goods (e.g. Cournot (1838), Buchanan and Yoon (2000), Feinberg and Kamien (2001)). I show that relaxing this assumption, allowing for degrees of complementarity, can alter the equilibrium outcome in a Pareto improving way. Several papers, including Feinberg and Kamien (2001), suggest imperfect complementarities should be considered. Parisi and Depoorter (2003) and Parisi, Depoorter, and Schulz (2005) consider imperfect complementarities, but with neither the same motivation nor objective as my model. Cournot (1838) considers the possibility of relaxing the perfect complementarities assumption, but he provides no analysis after concluding the model is too complicated to determine general results. In my consideration of imperfect complementarities and refunds in the assembly problem, I use a sequential model with

[2]Michelman (1982) defines *right* as "others are legally required to leave the object alone save as the owner may permit" and defines privilege as "the owner is legally free to do with the object as he or she wills."

discrete quantities, continuous price, two sellers, and one buyer, allowing me to derive enlightening results from a straightforward model.

2. General framework

Two monopolists, A and B, know a single buyer C finds value in combining component rights a and b to form composite right c.

For $J \in \{A, B\}$ and $j \in \{a, b\}$, seller J chooses price p_j at which to supply quantity $q_j \in \{0, 1\}$ of component j.[3] Let $p_c = p_a + p_b$ be the composite good's total price when components are supplied separately by monopolists A and B. While each seller J offers take-it-or-leave-it price p_j for component j, buyer C chooses to purchase ($q_j = 1$) or not purchase ($q_j = 0$) each component.

2.1. Information and timing

Players A, B, and C know all potential payoffs and all previous moves taken in the game. The order of play is fixed; seller A always moves first, and buyer C immediately follows each seller. A welfare analysis in Section 6 evaluates how order of play affects the anticommons outcome.

Fixed order of sequential play is not far-fetched. Geographical aspects of the market might naturally result in exogenously determined order. For example, a transportation system composed of multiple links joined by toll booths requires a buyer to purchase a pass at each link in order to proceed towards the destination. The composite good is the overall right to reach the destination. A fisherman might have his sights set on fishing from a lake high in the mountains. Such rights-seeker must obtain not only a permit to fish, but also permits to traverse private and/or government owned property. Buyers and sellers know a fishing license first must be obtained, and the geographical nature of getting to the fishing hole results in the exogenously

[3]Later I consider $q_a \in \{0, 1\}$ but $q_b \in \{0, 1, 2\}$.

determined process of obtaining the remaining permits to create the composite good, the right to fish at that lake.

The model rests on the assumption that sellers have the power in trade, where a take-it-or-leave it price is offered, and the buyer only can accept or reject an offer. A model of sellers' take-it-or-leave-it offers allows for examination under the most extreme outcomes of anticommons. This offers the most robust analysis of the role of refunds and stand alone value in mitigating hold-up.

If the buyer had the power of trade, rather than sellers, then trade occurs only if the buyer knows the sellers' valuation of their components. In this case, efficient trade results, and clearly, hold-up cannot occur.

A third mechanism of trade is when neither party has the power in trade and bargaining is possible. Under a bargaining model with fully transparent information and perfect complementarities, an outcome of hold-up depends on the split of cooperative surplus between the seller and the buyer at each stage of bargaining. Even under bargaining, an assumption of rationality results in all players accounting for sunk cost, thereby disregarding the purchase price of component a in the bargain between seller B and buyer C. Even if buyer C captures some of the gains to trade in the bargain with B, if his surplus from his bargain with A yields a negative surplus in sum, then hold-up results. However, if buyer C's bargaining power at both bargaining stages is high enough, hold-up might not result, due only to the buyer's bargaining power, not necessarily other factors such as refunds and stand alone value.

Therefore, a bargaining model is more involved than a model of take-it-or-leave it offers because outcomes depend on the interaction of bargaining power with refunds and/or stand alone value. The goal of this chapter is to show how refunds and stand-alone-value overcome outcomes of anticommons. A sequential model of perfect and complete information with sellers offering take-it-or-leave-it prices provides the best framework with which to highlight the impact of refunds and stand alone value.

2.2. Parameters and assumptions

Next, I explicitly define and distinguish the model's exogenous parameters, refunds and complementarities. Refunds alter the market structure for composite c, while complementarities are inherent to the production technology of a component. My analysis relies on varying the values of these parameters.

Any seller other than the last may offer a refund. Therefore, let $\gamma \in [0, 1]$ be the fraction of purchase price p_a seller A offers as a refund to buyer C. Regardless of whether components are complementary, a refund may be offered.

Definition 1 (refund)**.** ***No refund*** *corresponds to level* $\gamma = 0$*,* ***partial refund*** *corresponds to level* $0 < \gamma < 1$*, while a* ***full refund*** *corresponds to level* $\gamma = 1$*.*

The degree of complementarity is expressed by buyer C's willingness-to-pay for stand alone units of each component. Buyer C's willingness-to-pay for composite c and stand alone units of a and b is $\omega_c > 0$, $\omega_a \geq 0$ and $\omega_b \geq 0$, respectively. I assume buyer C's willingness-to-pay for the composite good is higher than the willingness-to-pay for both component goods separately. This focuses the analysis to that of a fragmented market for composite c.

Assumption 1 (Preference for composite)**.** *Buyer C prefers composite c to components a and b separately:* $\omega_c > (\omega_a + \omega_b) \geq 0$*.*

A component might have value only in use with other components to form the composite good. Suppose component a has value to buyer C only to form composite c. In this case, component a is a perfect complement to component b, even though b may or may not be perfect complements with a.

Definition 2 (perfectly complementary)**.** *The use for composite c is* ***perfectly complementary*** *if all components required to form the composite have no value to the buyer in use alone.*

A component might have value not only in use to create composite c, but also in use alone. For example, although component a is required to form composite c, component a also might have outside value to the buyer. In this case, even if the buyer cannot form composite good c, an outside use for component a exists.

Definition 3 (stand alone value, imperfectly complementary)**.** *For $j \in \{a, b\}$, a component has value beyond its use to create the composite, or* ***stand alone value (SAV$_j$)****, if $\omega_j > 0$. The use for composite c is* ***imperfectly complementary*** *if at least one component j has stand alone value.*

With two sellers and one representative buyer, the dynamic game of perfect and complete information is composed of four stages, generally outlined in Figure 1.1.

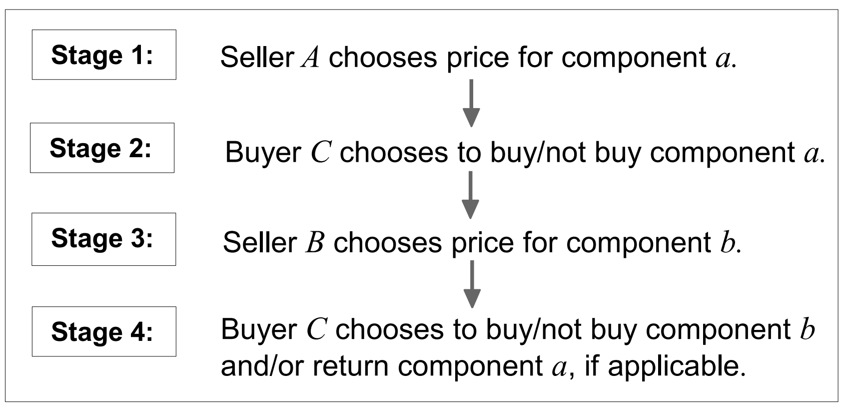

Figure 1.1. Game Sequence

My analysis uses the subgame perfect equilibrium concept, solving via backward induction. A subgame perfect equilibrium always exists because the game is finite with perfect information. Some games in my model result in multiple equilibria. However, for each game, the outcome is unique with respect to whether an anticommons tragedy results.

I make two equilibrium selection assumptions regarding indifference. These assumptions focus the equilibrium outcomes on those concerning an anticommons tragedy.

First, each seller, A and B, incurs a positive cost to supplying any positive quantity of their respective component. These costs could be production costs if components are physical goods, or transaction costs to trading rights. In either case, I call these costs transaction costs.[4] This assumption leaves seller A not indifferent when optimally choosing prices.

Assumption 2 (Transaction Costs)**.** *The cost of supplying any positive quantity of component j is $\epsilon_j > 0$. Furthermore, (i) if either component has positive stand alone value, $\omega_j > 0$, assume $\omega_j > \epsilon_j > 0$ and (ii) if either component has no stand alone value, $\omega_j = 0$, assume $\omega_c - \omega_{\not{j}} > \epsilon_j > 0$ for each $j \in \{a, b\}$.*

Assumption 2 ensures that any outcome in which the composite good is not purchased is due to the anticommons tragedy and not high transaction costs. When transaction costs are positive, each seller must ensure his optimal price yields non-negative payoffs, given the optimal behavior of the other two players. Thus, each seller must account for his positive transaction costs. Parts (i) and (ii) of Assumption 2 provide additional requirements on transaction costs to ensure supplying the component is optimal for a seller.

Assumption 2 part (i) requires that supplying one unit costs less than the buyer's willingness-to-pay for a stand alone unit of that component. This ensures that supplying a component is optimal whenever a component has positive stand alone value. However, when monopolist J's component has no stand alone value, Assumption 2 part (ii) ensures monopolist J optimally supplies his component, j, if the stand alone value of the other component, $\not{j}$, is not too high. This is because the degree of stand alone value of the other component, $\not{j}$, might determine how much surplus supplier J must leave on the table to ensure the buyer purchases his component, j.[5]

[4]For simplicity, I assume fixed costs rather than marginal costs. Given that most of the analysis restricts capacity to exact proportions, fixed and marginal costs of supplying that unit are equivalent. Once I remove the capacity constraint in Section 4.2.1, fixed costs still simplify the analysis; however, I indicate how marginal costs alter the outcome, requiring marginal cost be less than marginal benefit.

[5]Proofs of each Lemma, Proposition and Corollary, provided in Section 8: Appendix A, state the

Additionally, I assume if the buyer is left indifferent between more than one action, then the indifferent buyer chooses the action that results in the greatest purchased quantity of the component.

Assumption 3 (Indifference Rule)**.** *An indifferent buyer chooses to purchase the maximum possible units of the good.*

In sum, in this two-seller, one-buyer model, composite c is formed from one of the following three structures of components a and b: (i) two perfectly complementary components; (ii) two imperfectly complementary components; or (iii) a hybrid of complementarities. This amounts to four cases:

1. perfect complementarities market for c: $\omega_a = 0$, $\omega_b = 0$ (Baseline case);
2. stand alone value in only component a (SAV_a): $\omega_a > 0$, $\omega_b = 0$;
3. stand alone value in only component b (SAV_b): $\omega_b > 0$, $\omega_a = 0$; and
4. stand alone value in both components a and b ($SAV_{a,b}$): $\omega_a > 0$, $\omega_b > 0$.

Each case is analyzed under the three levels of refund (no, partial, and full refunds), leaving a total of twelve cases I consider.

A complete description of the game includes the strategy set for each player. Accounting for refundability, the strategies for each player at each decision node are given by:

$$S_A = p_a \geq 0$$

$$S_B = \{p_{b0} \geq 0, p_{b1} \geq 0\}$$

$$S_C = \{q_a \in \{0,1\}, q_{b0} \in \{0,1\}, q_{b1} \in \{0, 0_{\not a}, 1_c, 1_{a,b}, 1_{\not a,b}\}\}$$

For buyer C's strategy S_C, the first parenthesized term represents action to not buy/buy a; the second parenthesized terms represents action to not buy/buy b if

exact Assumption 2 conditions required in solving for the equilibrium.

$q_a = 0$; and the third parenthesized terms represents action to not buy/buy b with an option to return a (noted $\not a$) when $q_a = 1$. Figure 1.2 provides a description of the notation used throughout my analysis and in the extensive form games.

Payoffs:
Seller A
Seller B
Buyer C

p_j = price of j = action for J, $j \in \{a, b_0, b_1\}$, $J \in \{A,B\}$
q_j = quantity of j purchased = actions for C, $j \in \{a, b_0, b_1\}$
ε_j = transaction cost for seller J, $j \in \{a,b\}$, $J \in \{A,B\}$
$0 \leq \gamma \leq 1$: fraction of purchase price p_a refunded
δ = additional cost for A (surplus left on table for B) when $\gamma = 1$

$0_{\not a}$: don't buy b and return a
$1_{\not a,b}$: buy b but return a
$1_{a,b}$: buy b; use a and b separately
1_c : buy b; join a and b to create c
$2_{c,b}$: buy b; use one b to create c, use the other b separately
$2_{b,b}$: buy b; use both b separately

SAV_j: component(s) j have positive stand alone value, $j \in \{a,b\}$

FIGURE 1.2. Legend for notation

Figure 1.3 exhibits the most general form of the game. This general form of the game is simpler than it appears. First, notice the game simplifies significantly when the refund is zero. Even with a positive refund, several of buyer C's stage 4 actions are never played in equilibrium. For example, for any game considered in this analysis, by Assumption 1 (preference for the composite good) buyer C never finds it optimal to purchase both components to use separately ($q_{b1} = 1_{a,b}$) over purchasing both components to form composite c ($q_{b1} = 1_c$).

Finally, in the general form game, consider the second-smallest subgame along the path $q_a = 0$. In this $q_a = 0$-subgame, seller B chooses p_{b0}, then buyer C chooses q_{b0}. If component b has stand alone value, as long as supplying b has the lowest opportunity cost, seller B prices to sell stand alone units of b at the monopoly price. Buyer C

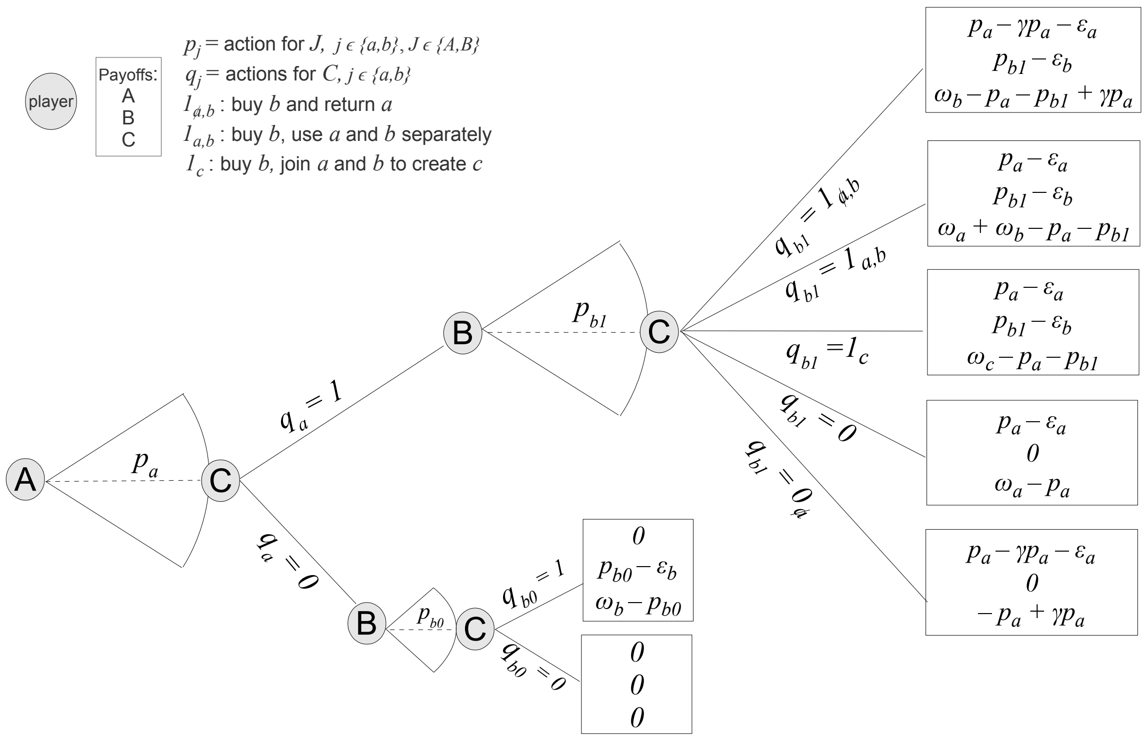

FIGURE 1.3. Game of refund $\gamma \in [0, 1]$ with nonnegative SAV_a and SAV_b

then purchases stand alone units of b at the monopoly price, and seller A receives zero payoff. If component b does not have stand alone value, then by Assumption 2, B prices above zero, resulting in no units of b purchased by buyer C in the subgame equilibrium.

This Nash equilibrium is off the subgame perfect equilibrium path when no anti-commons tragedy results and along the equilibrium path when hold-up occurs. However, the focus of the analysis is on averting hold-up. While the equilibrium outcome of each game depends on optimal behavior along the $q_a = 0$ path, regardless of the parameter values, the buyer always compares an equilibrium payoff of zero along the $q_a = 0$ path to the anticipated payoff along the $q_a = 1$ path. Therefore, in referencing the extensive form game tree for each game, focus should remain on the $q_a = 1$ path, where the tragedy ultimately might be averted. Lemma 1 summarizes the equilibrium along this path that is realized only when hold-up results.

Lemma 1. ($q_a = 0$ subgame) *By Assumptions 2 and 3, the Nash equilibrium of the second-smallest subgame along path $q_a = 0$ of the most general form of the game is*

$$\text{(i) if no } SAV_b,\ (p^*_{b0} > \epsilon_b, q^*_{b0} = 0) \text{ and (ii) if } SAV_b \text{ or } SAV_{ab},\ (p^*_{b0} = \omega_b, q^*_{b0} = 1).$$

In equilibrium, if $q_a = 0$, the final good is not purchased, but if component b has stand alone value, stand alone units of b are purchased.

The two most extreme cases of the general game are considered in Proposition 1, Proposition 6, and Proposition 7. The least restrictive case allows for positive refunds and stand alone value in both components. The most restrictive scenario, which serves as my baseline case, assumes no refunds with perfect complements. Feinberg and Kamien (2001) use a toll booth example to show the hold-up outcome in this most restrictive case.

3. Baseline: Game of no refund and perfect complementarities

For a concrete example, consider the following scenario used throughout the analysis.

> **Example:** Two dilapidated buildings each sit on a quarter-acre parcel of land, a and b. Owners A and B were grandfathered into a city zoning ordinance that prevents the building of new dwellings on any parcel less than a half-acre. Upon sale of either quarter-acre plot, the grandfather clause ends, so a new owner may not construct a new dwelling on either parcel. A land developer desires to combine exactly these two parcels to create a legal-size parcel to build a new home, compliant with the zoning ordinance. No substitute parcels exist. The developer values the combined right to both parcels more than the right to using each property separately. Owner A first offers the developer a take-it-or-leave it price.

> Upon observing the price, the developer decides whether to purchase parcel a. All parties know the developer has no value for either property (building or land) in use alone. Also, the purchase of parcel a does not include any form of earnest money agreement, options agreement, or other contract regarding refunds. Once all parties observe whether the developer purchases parcel a, Owner B offers a take-it-or-leave it price for parcel b. Finally, the developer chooses whether to purchase b.

Assume no refund, $\gamma = 0$. Under perfect complementarities, Buyer C demands each component only when combined to form composite c. The buyer's reservation price for using either component alone is zero, $\omega_j = 0 \; \forall j \in \{a, b\}$.

> **Example continued:** Due to the zoning restrictions, the developer cannot build a new home unless he acquires both parcels of land. Because the developer finds value only in building a new home and has no value for either parcel in use alone, buildings a and b are perfect complements to the developer.

The extensive form game of perfect complements and no refund, given in Figure 1.6, shows the simplified general game. Without the option to refund component a, Buyer C's stage 4 actions are limited to choosing not to purchase b, to purchase b to combine with a to form c, or to purchase b for use alone. Without stand alone value, clearly b for stand alone use is never chosen.

Proposition 1. (No Refund/Perfect Complements - Tragedy) *Assume seller A offers no refund. By Assumptions 2 and 3, if components a and b are perfect complements, then in equilibrium, buyer C does not purchase either component good. Therefore, in equilibrium, buyer C does not purchase the composite good.*

Example continued: Because parcel a has no use except when combined with parcel b, Owner B, the last seller, knows if the developer purchases parcel a, the purchase price of a is fully sunk. Therefore, b rationally charges the developer's willingness-to-pay for the combined parcels. Given the purchase of parcel a, the developer rationally would purchase parcel b at this price. However, unless the purchase price for parcel a is zero, the developer knows his *ex post* surplus will be negative. If the cost to supplying parcel a is positive, then owner A will never price at zero. Therefore, the developer will never purchase parcels a or b.

In a sequential setting, the anticommons tragedy under perfect complements and no refund amounts to full hold-up. Once the buyer purchases from first seller A, the cost for component a is sunk. The sunk nature of this purchase reveals itself in the final payoffs to buyer C: once buyer C purchases a, for all of buyer C's stage 4 actions, C's payoff is reduced by the purchase price, p_a. This sunk cost created from the transaction between seller A and buyer C yields a positive externality on seller B. With the cost of a sunk, seller B optimally prices not as a monopolist for component b, but as a monopolist for composite good c. This results in negative buyer surplus for C for any $p_a > 0$. However, for positive transaction costs, pricing component a at zero yields a negative payoff to seller A. Therefore, seller A optimally prices above zero. In equilibrium, the total price exceeds buyer C's willingness to pay for composite c, so buyer C does not purchase a or b. Equilibrium surplus is zero for every player.

This discussion clarifies the critical role of sunk costs in causing the anticommons tragedy in a sequential model. I now consider how stand alone value and refund might eliminate the tragedy by reducing sunk costs.

4. Games of no refund and imperfect complements

I now alter the baseline of perfect complements by allowing positive stand alone value in either one or both component goods. None of the players' action sets change, but with changes in the stand alone parameter, payoffs change.

4.1. No refund and SAV_a and SAV_{ab}

Begin by assuming $\omega_a > 0$ and $\omega_b \geq 0$ so *at least* component a has stand alone value to buyer C.

> **Example continued:** Suppose parcel a has stand alone value to the developer. For example, even though the developer cannot use only parcel a to build a new home, he can use the dilapidated building on parcel a as a storage unit, which has value to him, but not as much value as building a new home on parcels a and b combined. Now parcel a, the first parcel offered for sale, has stand alone value to the developer.

Figure 1.7 provides the merged, extensive form games of no refund with stand alone value in at least component a; the two possible games, SAV_{ab} and SAV_a differ only in the payoffs for actions $q_{b1} = 1_{a,b}$ (purchase b to use with a) and $q_{b0} = 1$ (purchase b for stand alone use only).

Whether the game is one of stand alone value in both components or stand alone value in a alone, Assumption 1 implies buyer C will never choose to purchase both components to use separately (action $q_{b1} = 1_{a,b}$). Therefore, the only difference in the equilibrium of these two games of stand alone value in a is in the stand alone market for component b. The outcome of both games, with respect to anticommons, is the same. For this reason, I combine the analysis of games of imperfect complementarities in component a.

Proposition 2. (No Refund/SAV_a or SAV_{ab} - No Tragedy) *Assume seller A offers no refund. By Assumptions 2 and 3, if at least first component a has stand alone value, then buyer C purchases the composite good in equilibrium.*

> **Example continued:** Owner B knows once parcel a is purchased, the cost is not fully sunk because parcel a has use to the developer other than in use with parcel b. Therefore, the developer will not be willing to pay the monopoly price for the the assembled property, c, given the purchase of parcel a. Owner B optimally offers a lower price to ensure purchasing b is optimal for the developer. Because b does not hold-up the developer, owner A can price positively and still ensure the developer purchases parcel a.

For both games, the subgame perfect equilibrium along the equilibrium path is

$$(p_a^* = \omega_a, (q_a^* = 1, q_{b1}^* = 1_c), p_{b1}^* = \omega_c - \omega_a).$$

Along the equilibrium path, the composite good is purchased. Only off-the-equilibrium-path equilibria differ between a game of SAV_{ab} and a game of SAV_a.

Proposition 2 offers two implications. First, in equilibrium, positive stand alone value in the first component is enough to prevent full hold-up by the last seller, regardless of whether a refund is offered or the last component has stand alone value. If the first component has stand alone value, the sunk cost to purchasing component a is *reduced* by ω_a, buyer C's stand alone value of a. This reduction in sunk cost via an increase in component a's stand alone value gives the last seller incentive to charge only the composite good monopoly price *less* this change in sunk cost (ω_a). Without incentive to hold-up, enough surplus is left on the table for seller A optimally to price positively and buyer C optimally to purchase both components.

Second, in equilibrium, both sellers earn positive profits, compared to zero profits realized in the perfectly complementary market. Additionally, if component b has stand alone value, then seller B charges $p^*_{b1} = \omega_c - \omega_a$, a higher equilibrium price for a unit of b to be combined with a than could be charged in the stand alone market for component b when buyer C does not purchase component a, $p^*_{b0} = \omega_b < \omega_c - \omega_a$. Therefore, last seller B benefits from the strategic interaction with seller A in the market for composite c, as opposed to being the sole monopolist in the stand alone market for component b.

4.2. No Refund and SAV_b

> **Example continued:** Suppose only parcel b has stand alone value to the developer. For example, even though the developer cannot use parcel b alone to build a new home, he can rent the building on parcel b. However, the total value of renting the building on parcel b is lower than the value of building a home on assembled parcel c.

A game of no refund with stand alone value in component b assumes first component a has no stand alone value ($\omega_a = 0$), last component b has positive stand alone value ($\omega_b > 0$), and seller A offers no refund ($\gamma = 0$). Figure 1.8 provides the extensive form of the game. Notice that the game's payoffs differ from the perfect complements/no-refund case examined in Section 4 in only two outcomes: the case when buyer C chooses action $q_{b1} = 1_{a,b}$ (purchasing a and b to use alone) and the case when buyer C chooses action $q_{b0} = 1$ (purchasing b to use alone). Therefore, since rational buyer C never chooses $q_{b1} = 1_{a,b}$, equilibrium hold-up results, as in the game of perfect complements with no refund. Stand alone value in component b does not reduce the sunk cost to purchasing a. Optimal behavior yields hold-up, with the total price in equilibrium higher than the buyer's willingness to pay for composite c. Unlike the perfect complements case, since b has stand alone value, a stand alone unit of b is purchased in equilibrium.

Proposition 3. (No Refund/SAV_b - Tragedy) *Assume seller A offers no refund. By Assumptions 2 and 3, if last component b has stand alone value, then buyer C does not purchase the composite good but does purchase a stand alone unit of b in equilibrium.*

> **Example continued:** Owner B knows parcel a has no stand alone value to the developer, so the purchase of parcel a is fully sunk. Owner B optimally prices at the value to the developer for combining both parcels. In equilibrium, this holds up the developer because the total equilibrium price exceeds his willingness-to-pay for both a and b. The developer does not purchase a. Since he values parcel b alone, and since B observes that the developer has not purchased a, owner B offers a price equal to the developer's value for b. The developer therefore purchases parcel b at the price he is willing to pay for using b alone as a rental property.

Some components, such as real or intellectual property, as in the examples I have considered, have a natural capacity constraint of one unit. However, allowing the last supplier to sell more units of b than are required to create composite c can alter the optimal behavior in this game. To understand whether stand alone value in the last component supplied affects hold-up, I relax the capacity constraint on component b.

4.2.1. No price discrimination and relaxed capacity constraint

> **Example continued:** Now suppose Owner B offers for sale a second parcel of similar land having the same stand alone value on the opposite end of owner A's parcel. Either of the b-parcels could be combined with parcel a to form the larger parcel, with the extra b-parcel used alone.

Suppose in stage 4, buyer C has the option to purchase two units of component b, with the quantity of component a supplied held constant at one unit. Therefore, seller B knows if buyer C purchases component a and values b alone, then the two units of component b purchased must be for different uses – one for creating c and one in use alone. If price discrimination were feasible, B would offer different prices for the first and second units of b. If B cannot price discriminate, B must choose a single, per-unit price for component b.

If buyer C purchases component a, then C's set of actions in stage 4 becomes (i) purchase zero units of b ($q_{b1} = 0$); (ii) purchase one unit of b to form composite good c ($q_{b1} = 1_c$); (iii) purchase two units of b: one to form good c and one to use alone ($q_{b1} = 2_{c,b}$); (iv) purchase one unit of b to use alone ($q_{b1} = 1_{a,b}$); and (v) purchase two units of b to use each alone at value ω_b ($q_{b1} = 2_{b,b}$).

One can easily show buyer C never chooses actions (iv) or (v). However, under certain conditions, buyer C might choose to purchase two units of b: one to form c and the other to use b alone. If buyer C does use one unit of b alone, he gets value ω_b for that unit. Therefore, if C purchases two units, he will never agree to pay more than the value of b used alone. As a result, seller B knows he cannot charge a price greater than C's willingness-to-pay for b alone.

However, selling two units of b at a per-unit price no higher than ω_b might not yield higher payoffs to B than selling just one unit of B at a price higher than ω_b. Seller B knows C would purchase one unit of b to use with a at a price no higher than ω_c. In that case, B could price above ω_b and no higher than ω_c. Then, seller B earns profit for selling only one unit of b, but at a price higher than ω_b. Therefore, seller B prices to sell either two units of b or one unit of b, depending on the relationship between ω_b and ω_c.

Lemma 2 formalizes this pricing incentive by analyzing the Nash equilibrium of the second smallest subgame in the full game. This Lemma highlights the importance of the degree of stand alone value in affecting optimal behavior. The following definition provides a measure of the degree of stand alone value:

Definition 4 (high (low) stand alone value)**.** *For $j \in \{a, b\}$, component j has* ***high stand alone value*** *to buyer C if $2\omega_j > \omega_c$. Otherwise, component j has* ***low stand alone value****.*

Thus, if a component's stand alone value is more than half of the value of the composite good, then that component has high stand alone value. Otherwise, that component has low stand alone value. If both components, a and b, have high stand alone value, then $2\omega_a + 2\omega_b > \omega_c + \omega_c \Leftrightarrow \omega_a + \omega_b > \omega_c$, which contradicts Assumption 1. Therefore, at most one component can have high stand alone value. If both components have low stand alone value, Assumption 1 is satisfied.

If component b has high stand alone value, seller B optimally prices to sell two units, even though he cannot price above ω_b in doing so. However, if component b

has low stand alone value, B optimally charges a higher price and sells only one unit of b. It is not optimal to price lower to sell two units when the stand alone value, and thus the feasible price, of the units is low.

Lemma 2. (No Refund/SAV_b unconstrained capacity: B's pricing incentives) *Assume component b has stand alone value. Consider the second smallest subgame along branch $q_a = 1$ in a game of no refund with stand alone value in last component b. Assume C may purchase up to two units of component b, and seller B cannot price discriminate. By Assumptions 2 and 3, (i) seller B optimally prices at $p_{b1}^* = \omega_b$ to sell $q_{b1}^* = 2_{c,b}$ if and only if last component b has high stand alone value and (ii) seller B optimally prices at the monopoly price for the composite good, $p_{b1}^* = \omega_c$, to sell $q_{b1}^* = 1_c$ if and only if last component b has low stand alone value.*

Figure 1.4 illustrates how the demand for component b is affected by the relationship between the stand alone value of component b and the reservation value of composite c. The demand for a component b with high stand alone value, ω_{b-high} in the graph, is more elastic than demand for a component b with low stand alone value, ω_{b-low} for the following reasons. By Assumption 1, since buyer C values the composite good to using either component alone, then regardless of the stand alone value in b, the buyer is willing to pay up to ω_c for the first unit of b. For the second unit of b, buyer C only will pay up to his stand alone value for the component, either ω_{b-low} or ω_{b-high}. In the standard monopoly model, the monopoly price under more price elastic demand is lower than the monopoly price under a less price elastic demand. This is exactly the result shown in Lemma 2; by Assumption 1, with high stand alone value in b, seller B's optimal price (ω_b) is lower than the monopoly price for composite c (ω_c), which is B's optimal price when b has low stand alone value.

Using B's optimal pricing strategy given in Lemma 2, Proposition 4 provides conditions for which buyer C purchases composite c, yielding no anticommons tragedy, and conditions for which B holds up C, resulting in an anticommons tragedy.

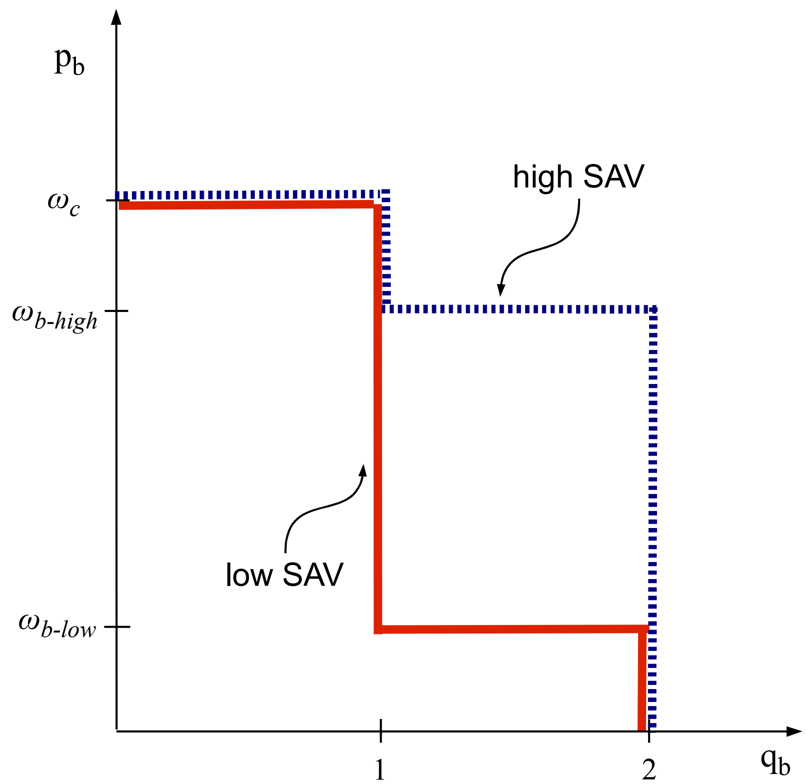

Figure 1.4. Demands for component b by degree of stand alone value

Proposition 4. (No Refund/High(Low) SAV_b - No Tragedy (Tragedy)) *Assume seller A offers no refund and only last component b has stand alone value. Assume buyer C may purchase up to two units of component b, and seller B does not price discriminate. By Assumptions 2 and 3, in equilibrium, (i) if last component b has high stand alone value, then the buyer purchases the composite good and (ii) if last component b has low stand alone value, then the buyer does not purchase the composite good, but does purchase stand alone units of b.*

Compared to the baseline of perfect complements, if buyer C purchases component a, the sunk cost still equals the purchase price of component a, p_a. However, with the option of selling more than one unit of b, the last seller can increase profits by reducing the price to sell two units of b. Thus, when buyer C highly values component b, seller B has an incentive not to hold-up. Therefore, stand alone value in the last component b can be enough to mitigate the anticommons tragedy when capacity is not constrained to the exact proportion required to create the composite good.

4.2.2. Price discrimination and relaxed capacity constraint

Once the last seller offers to sell more than one unit of component b, it is natural to suppose seller B will price discriminate. If buyer C purchases more units of a component than required to form the composite good, then the seller knows the buyer intends to use those additional units for an alternate, less valuable use. By Assumption 1, seller B knows buyer C places value $\omega_c > \omega_b$ for any units of b purchased to create composite c. Therefore, buyer C is willing to pay up to ω_c for these units. The buyer will pay no more than reservation price ω_b for any additional units of b purchased. Thus, seller B might first degree price discriminate by charging a price $p_{b1c} \leq \omega_c$ for any units of b used to create c and price $p_{b1b} \leq \omega_b$ for any units of b not used to create composite c. Clearly, B prices to capture all buyer surplus through perfect price discrimination because reservation prices are known. This provides no incentive for A to price so that buyer C purchases component a. Therefore, the composite is not purchased in equilibrium.

Proposition 5. (No Refund/SAV_b: Price Discrimination - Tragedy) *Assume seller A offers no refund and component b has stand alone value. Assume buyer C has the* *option to purchase up to two units of component b. By Assumptions 1, 2, and 3, if seller B perfectly price discriminates, then in equilibrium, buyer C does not purchase composite c but does purchase two stand alone units of b.*

Thus, hold-up results if the last seller price discriminates. This inefficient outcome depends on an assumption of perfect and complete information. In real world examples of assembly problems, information can be transparent due to close relationships between parties. In such cases, sellers can determine a buyer's willingness-to-pay, so allowing for price discrimination can result in the anticommons tragedy.

Price discrimination by seller B imposes a negative externality on seller A because it leaves A without an optimal pricing strategy that gives buyer C incentive to purchase component a. With or without price discrimination, seller B sells two units of component b. However, with price discrimination, the composite good is not purchased in equilibrium. Since buyer C values stand alone units of b, seller B is indifferent to price discriminating or not. However, when B price discriminates, seller A is worse off, and the anticommons tragedy results. Recall, if component b has high stand alone value and B does not price discriminate, hold-up is averted. Therefore, when b has high stand alone value, banning price discrimination yields a weak Pareto improvement.

Figure 1.5 summarizes the results from games with stand alone value in component b. Whenever A does not offer a refund and b has stand alone value, then the anticommons tragedy is mitigated if and only if (1) C has the option to purchase more than one unit of b, *and* (2) b has high stand alone value, *and* (3) B does not engage in price discrimination. If C cannot purchase more than one unit of b, or if b has low stand alone value, or if B price discriminates, then B prices to capture all of the surplus from buyer C, leaving seller A unable to profitably sell any units of a. In

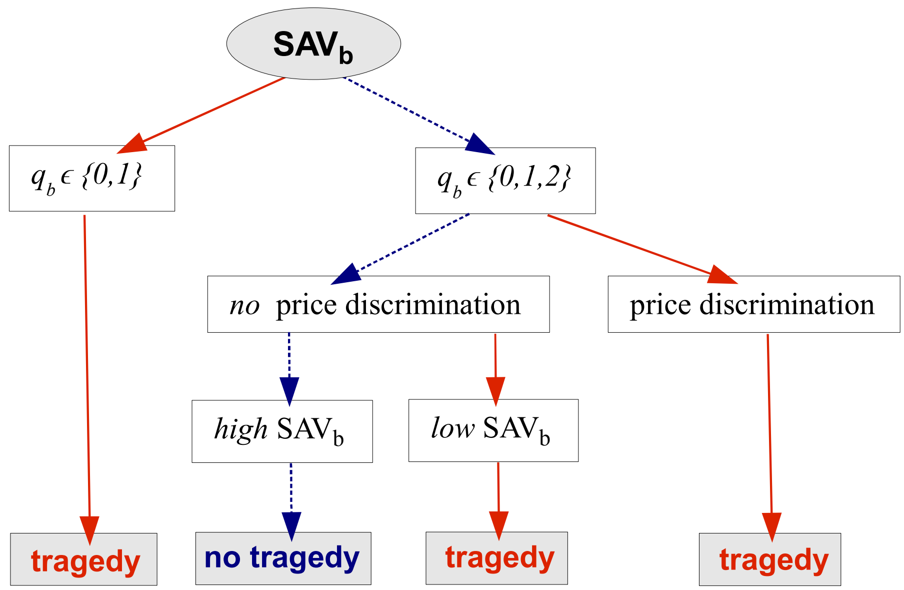

FIGURE 1.5. Results: Games of stand alone value in only component b

such cases, C does not purchase a and does purchase stand alone units of b, resulting in an anticommons tragedy.

5. Games of refunds

I now incorporate positive refundability, $\gamma \in (0, 1]$. I consider partial refund, $\gamma \in (0, 1)$, and full refund, $\gamma = 1$.

> **Example continued:** Now, assume a contract between owner A and the developer includes an earnest money option for the purchase of parcel a. Rather than outright purchasing parcel a, the developer and owner A agree upon a value of earnest money to be paid to seller a, enough to signal the developer's intent to purchase, but not too much to dissuade the developer from participating. Only one unit of parcel b is available for

purchase. Under the agreement, if the developer does not purchase parcel b, some portion of the earnest money for parcel a is refunded and parcel a is not sold.

As shown above, B's ability to hold-up the buyer derives from the sunk cost of purchasing component a. A full refund completely eliminates this sunk cost, while a partial refund reduces it. Figure 1.10 again provides the extensive form of the general game, allowing for all degrees of complementarity and refunds. Notice how refunds reduce sunk cost, either fully or partially, when C chooses to return a.

Consider the perfect complements case. Since component a does not have stand alone value, C's payoff will be negative in the event he does not purchase component b. When seller A offers a partial refund, full recovery of the purchase price for a does not occur. However, a full refund makes buyer C no worse off when purchasing a. Therefore, a partial refund, which leaves the buyer without a potential for positive surplus, is not enough to change seller B's incentive to hold-up. On the other hand, with a threat of full refund, B must price low enough to ensure C has incentive to purchase both components. In fact, with sunk costs fully eliminated, B's best response directly depends on seller A's optimal behavior. A captures (nearly) all surplus, ω_c, leaving just enough surplus on the table to ensure an indifferent seller B optimally prices to sell component b.

Let $\delta > 0$ be this added surplus seller A leaves on the table for B to optimally price so C chooses $q_b = 1_c$ and to ensure B is not indifferent. Therefore, this smallest element of surplus acts as a cost to seller A to ensure buyer C purchase component a in equilibrium.

Proposition 6, below, summarizes the partial refund's weakness as a mechanism for mitigating hold-up. If buyer C purchases the first component but not the last, then the buyer is better off *ex post* with even partial recovery of the sunk cost. However, this partial recovery is not enough to alter the optimal behavior of seller B, so *ex ante*, the buyer has no incentive to purchase the first component anyway. This result holds regardless of the degree of complementarities. Therefore, under cases of complementarities and no refund where hold-up results, a partial refund will not avert an anticommons tragedy. Since partial refunds do nothing to alter the optimal behavior of B, when at least component a has stand alone value, full hold-up is mitigated irrespective of a partial refund. Even so, I show below that seller A might prefer a partial refund to a full refund. Therefore, a partial refund might not affect the anticommons outcome, but it can affect the equilibrium welfare.

Proposition 6. (Partial Refund) *Assume seller A offers only a partial refund $\gamma \in (0, 1)$. Composite c is purchased in equilibrium if at least component a has stand alone value. (Therefore, composite c is not purchased in equilibrium if components a and b are perfect complements or if only component b has stand alone value.) If component b has stand alone value, the stand alone unit of b is purchased in equilibrium.*

Though partial refunds do nothing to mitigate the anticommons tragedy, full refunds do. Proposition 7 shows conditions for each case of complementarities that ensure hold-up is mitigated when seller A offers a full refund. The result relies on seller A leaving just enough surplus on the table, δ, to induce seller B to sell component b. As long as this additional transaction cost to seller A, δ, is low enough, hold-up is mitigated.

Proposition 7. (Full Refund) *Assume seller A offers full refund, $\gamma = 1$. (i) If components are perfect complements or if only component a has stand alone value, the composite good is purchased in equilibrium as long as $\omega_c - \epsilon_a > \delta > \epsilon_b$. (ii)*

If at least component b has stand alone value, the composite good is purchased in equilibrium as long as long as $\omega_c - \omega_b - \epsilon_a > \delta$ *holds.*

With the option to return component a for a full refund, seller A must price to ensure seller B does not price so component a is returned: A wants B to price so b is purchased and C keeps a. When component b does not have stand alone value, assuming transaction costs are not too high, the condition ensuring these optimal actions is weak.

However, the option to return component a for a full refund increases buyer C's strategy set in a way that could benefit B. In the case where at least component b has stand alone value, if A prices "too high," C's optimal action may be to return a and consume b alone. In this case, to ensure seller B prices so buyer C will not return a, the stand alone value of component b (plus a small δ) must be left as surplus for seller B. Therefore, as shown in Table 1.1, when at least component b has stand alone value, a refund option can yield seller B significantly nonzero ($\omega_b + \delta$) equilibrium payoffs.

6. Welfare analysis

Using the results from above, outlined in Table 1.1, I consider how welfare is affected by degree of complementarity, refunds and the order of play. Endogenizing the choice of refund and endogenizing the order of play offer additional insight regarding optimal welfare.

6.1. Endogenous refund

Suppose seller A chooses a refund level in stage 0, and component a has stand alone value to buyer C. If the marginal gain to offering a full refund outweighs the marginal cost to offering a full refund, then seller A optimally chooses a full refund. This means C's stand alone value of a cannot be too similar to C's value for composite c, or in the case of stand alone value in only component a, B's transaction costs are not too high. Corollary 1 summarizes this result.

Corollary 1. (A chooses full refund) *In stage 0, let refund level* $\gamma \in [0, 1]$ *be chosen by seller A. A full refund* $\gamma^* = 1$ *maximizes seller A's payoffs in the cases where (i) component a has stand alone value to buyer C and* $\omega_c - \omega_a > \delta > \epsilon_b$ *and (ii) both components a and b have stand alone value to buyer C and* $\omega_c - (\omega_a + \omega_b) > \delta$.

Thus, when only component a has stand alone value to C, if a's stand alone value is too high, it might be optimal for seller A to relinquish payoffs to seller B by offering no or partial refund. A full refund is not necessarily optimal for the first seller.

As a matter of efficiency and ensuring no hold-up, a social planner should exogenously set a full refund in the case where the order of acquiring rights is fixed and either (i) only the last component required to complete the assembly of rights has stand alone value, or (ii) no components have stand alone value. This ensures full hold-up does not result. However, as a matter of wealth distribution, in the case where at least the first component has stand alone value, an exogenously set full refund could transfer surplus from the first seller to the last seller, if component a has high enough stand alone value or if B's transaction costs are high enough. Notice when only component b has stand alone value or no component has stand alone value, seller A always chooses a full refund because equilibrium payoffs are zero, otherwise.

6.2. Endogenous order

Next, relax the assumption of exogenously determined order of play. Below, the preferred order of play is considered in the cases when (i) the buyer chooses the order of play and (ii) sellers choose order of play.

6.2.1. Buyer chooses order

Buyer C's equilibrium surplus is zero in every case considered, above. Recall, by Assumption 1, buyer C participates in the market for composite c because C values the composite good more than using the component goods alone. Therefore, buyer C optimally chooses the order of play that yields no hold-up in equilibrium. Regardless of the level of refund, buyer C is guaranteed no hold-up as long as a component with stand alone value is purchased first. On the other hand, if a full refund is offered, the order in which buyer C purchases the components does not affect his ability to successfully create composite c in equilibrium.

6.2.2. Sellers choose order

A seller's equilibrium payoffs may be higher under a particular order of play. In this case, I consider when seller A has a first-mover advantage and when seller B has a second-mover advantage. Under no or partial refund, the degree of component a's stand alone value determines whether seller A has a first-mover advantage or seller B has a second-mover advantage.

Corollary 2. (First (Second) mover advantage) *Assume seller A offers no or partial refund. (i) If first component a has high (low) stand alone value, then $p_a^* > (<) p_{b1}^*$. (ii) If $\epsilon_a = \epsilon_b$ then $\pi_A^* > (<) \pi_B^*$.*

> **Example continued:** The developer's realized value in parcel a alone is not only enough to ensure Owner B does not hold-up, but if using a as a storage unit is not too valuable to the developer (use as a storage unit has low stand alone value), then in equilibrium Owner B could receive higher payoff than owner A. If owner A's building is not too damaged and the developer highly values a's building as a storage unit, then in equilibrium a could realize higher payoff than Owner B.

Regardless of whether b has stand alone value and no refund is offered, if component a has high stand alone value to buyer C, seller A's equilibrium price is higher than B's equilibrium price. A sufficient condition for a first-seller advantage is that transaction costs are equivalent and a has high stand alone value. A necessary condition for a first seller advantage is A's costs not be too high, compared to the stand alone value for a, or a's stand alone value be high enough, compared to A's transaction costs.

By similar reasoning and assumptions, under low stand alone value in component a, seller B may have a second-mover advantage. Therefore, even if C values use of a alone and does not value use of b alone, seller B still may earn higher equilibrium payoffs than seller A. Though stand alone value in a is sufficient to prevent hold-up by seller B, the degree of stand alone value determines which seller achieves higher payoffs in equilibrium. This result shows, for example, in a private takings situation, the homeowner waiting the longest to sell does not necessarily earn as high of payoffs than the first seller.

Corollary 3 shows that under a full refund, seller B could have a second-mover advantage. For stand alone value in at least component b, if that stand alone value is high, then seller B's equilibrium payoffs are higher than seller A's equilibrium payoffs.

Corollary 3. (Full refund second-mover advantage) *Assume seller A offers a full refund. If at least last component b has high stand alone value and* $\delta > \epsilon_b$ *(or* $\delta - \epsilon_b$ *is not too negative), then* $\pi_B^* > \pi_A^*$.

7. Summary of results and conclusion

Table 1.1 adds to the results given in Figure 1.5 and summarizes the equilibrium outcomes and equilibrium payoffs from the eight cases of complementarities and refundability, as well as the three cases relaxing capacity.

First, regardless of degree of complementarity, a full refund ensures the composite good is purchased in equilibrium. Therefore, a fragmented market of perfect complementarities does not result in hold-up if full refunds are allowed. Although full hold-up does not always result under partial refund, comparison of partial refund to no refund reveals the partial refund does nothing to mitigate the anticommons tragedy; the cases of no tragedy under partial refund are due to the stand alone value of component a. Tragedy does not result under no refund with at least stand alone value in the first component. When the first component does not have stand alone value, with a partial or no refund, tragedy results.

A comparison of equilibrium payoffs reveals, as expected in a game of completely transparent information, in equilibrium, the sellers fully capture the buyer's value of the composite good. Under the assumption an indifferent buyer purchases, then the buyer always is left with no surplus in equilibrium, but in cases of no tragedy, the composite good successfully is assembled.

7.1. Conclusion

Building on the general framework of complementary monopoly and anticommons property, this research first models the assembly problem in terms of sunk cost. Motivating the potential inefficiency in these fragmented markets through sunk cost reveals that solutions to hold-up can rely on factors that reduce those sunk costs. I show the effects of imperfect complementarities and refundability, two factors that reduce sunk cost, in mitigating hold-up. These two factors are relevant in real world examples of anticommons property and complementary monopoly. My analysis offers settings where stand alone value and a refund option can alter potential hold-up.

The order of purchase also affects the incentives for hold-up. Since hold-up is prevented when at least the first component purchased has stand alone value, a social planner should take notice at any required ordering to obtaining rights of use, efficiently determining order or allowing buyers to determine order based on their stand alone values for components.

When only the last component has stand alone value, hold-up can be mitigated if capacity is not constrained to only the proportion required to create the composite good. Allowing the buyer to purchase more units of the last component than required for assembly alters optimal behavior. However, in this case, the level of that last component's stand alone value relative to the composite good's value determines whether the last seller continues to have incentive to hold-up the buyer. This result implies that even though it appears stand alone value in the last component only heightens the last seller's incentive to hold-up, in fact, if additional stand alone units of the last component are supplied, the last seller may no longer have incentive to cause an anticommons tragedy.

Additionally, perfect price discrimination in these market settings results in an inefficient outcome. Because the last seller is indifferent to price discriminating, a ban on price discrimination is Pareto improving when the first component has high stand alone value.

Finally, full refunds mitigate the anticommons tragedy in my model, while a partial refund is no better than no refund for overcoming the anticommons tragedy. My welfare analysis reveals conditions for which a first seller or a last seller advantage exists; if a seller's component has stand alone value, that seller does not necessarily earn higher equilibrium payoff than his rival. Therefore, stand alone value in a rival's component could be a positive externality for the other seller.

Since the literature has not modeled refundability and a little research exists regarding imperfect complementarities, my analysis builds on the current research of anticommons property and complementary monopoly. Furthermore, the results suggest policies, such as legal requirements on full refunds, regulation on order in which components must be purchased, or legal prohibition of price discrimination, to mitigate inefficient outcomes in assembly problems.

Extensions to my research include a continuous, sequential market model of the assembly problem. I have investigated the continuous model under perfect complementarities with repeated, sequential play to gain insight into how repetition affects the incentive to hold-up; I hope to incorporate complementarities and refundability into such model. Finally, the model presented in this research lends itself to experimental testing in the laboratory; current research includes designing this experimental model. Also, I have considered a setting in which to empirically examine hold-up using naturally-occurring data, and I continue to determine the appropriate model and data requirements.

Refund level	**Case** PD = Price Discrimination	**Tragedy**	π_A^*	π_B^*	π_C^*
No Refund	Perfect Complements	yes	0	0	0
	SAV_a	no	$\omega_a - \epsilon_a$	$\omega_c - \omega_a - \epsilon_b$	0
	SAV_{ab}	no	$\omega_a - \epsilon_a$	$\omega_c - \omega_a - \epsilon_b$	0
	one unit SAV_b	yes	0	$\omega_b - \epsilon_b$	0
	two units SAV_b/PD	yes	0	$2\omega_b - \epsilon_b$	0
	two units Low SAV_b/No PD	yes	0	$2\omega_b^{\text{low}} - \epsilon_b$	0
	two units High SAV_b/No PD	no	$\omega_c - \omega_b - \epsilon_a$	$2\omega_b^{\text{high}} - \epsilon_b$	0
Partial Refund	Perfect Complements	yes	0	0	0
	SAV_a	no	$\omega_a - \epsilon_a$	$\omega_c - \omega_a - \epsilon_b$	0
	SAV_{ab}	no	$\omega_a - \epsilon_a$	$\omega_c - \omega_a - \epsilon_b$	0
	SAV_b	yes	0	$\omega_b - \epsilon_b$	0
Full Refund	Perfect Complements	no	$\omega_c - \delta - \epsilon_a$	$\delta - \epsilon_b$	0
	SAV_a	no	$\omega_c - \delta - \epsilon_a$	$\delta - \epsilon_b$	0
	SAV_{ab}	no	$\omega_c - \omega_b - \delta - \epsilon_a$	$\omega_b + \delta - \epsilon_b$	0
	SAV_b	no	$\omega_c - \omega_b - \delta - \epsilon_a$	$\omega_b + \delta - \epsilon_b$	0

TABLE 1.1. Summary of Results

8. Appendix A: Chapter 1 proofs

Proof of Lemma 1. Assume $q_a = 0$.

(i) Let $\omega_b = 0$. By Assumption 3, the indifference rule, buyer C chooses $q_{b0} = 1$ as long as $p_{b0} \leq 0$. Observing buyer C's optimal strategy, seller B chooses to solve

$$\max_{p_{b0}} \{p_{b0} - \epsilon_b, 0\} \text{ s.t. } p_{b0} \leq 0.$$

Therefore, the Nash equilibria are $\{p^*_{b0} > \epsilon_b, q^*_{b0} = 0\}$.

(ii) Let $\omega_b > 0$. By Assumption 3, the indifference rule, buyer C chooses $q_{b0} = 1$ as long as $p_{b0} \leq \omega_b$. Observing buyer C's optimal strategy, seller B chooses to solve

$$\max_{p_{b0}} \{p_{b0} - \epsilon_b, 0\} \text{ s.t. } p_{b0} \leq \omega_b.$$

Therefore, assuming transaction costs are low enough (Assumption 2(i) ($\epsilon_b < \omega_b$)), the Nash equilibrium is $\{p^*_{b0} = \omega_b, q^*_{b0} = 1\}$. □

Proof of Proposition 1. *(No Refund/Perfect Complements - Tragedy)*

Assume $\gamma = 0$, $\omega_a = 0$, and $\omega_b = 0$. In stage 4 and then stage 3:

(i) if $q_a = 0$, then by Lemma 1, the Nash equilibrium of the second smallest subgame is $\{p^*_{b0} > \epsilon_b, q^*_{b0} = 0\}$.

(ii) if $q_a = 1$, then in stage 4, by Assumption 1, buyer C never chooses $q_{b1} = 1_{a,b}$. Buyer C chooses $q_{b1} = 1_c$ if and only if

$$\pi_C(q_{b1} = 1_c) \geq \pi_C(q_{b1} = 0) \Leftrightarrow \omega_c - p_a - p_b \geq -p_a.$$

Therefore, by Assumption 3, the indifference rule, as long as $p_{b1} \leq \omega_c$, buyer C chooses

$q_{b1} = 1_c$. In stage 3, seller B solves

$$\max_{p_{b1}} \{p_{b1} - \epsilon_b, 0\} \text{ s.t. } p_{b1} \leq \omega_c.$$

Therefore, assuming transaction costs are low enough (by Assumption 2(ii), ($\omega_c > \omega_c - \omega_a > \epsilon_b$)), seller B chooses $p^*_{b1} = \omega_c$.

In the third smallest subgame at stage 2, buyer C, knowing the optimal actions of seller B in stage 3, chooses $q_a \in \{0, 1\}$. By Assumption 3, the indifference rule, buyer C chooses to purchase component a as long as

$$\pi_C(q_a = 1) = \omega_c - p_a - p^*_{b1} = -p_a \geq \pi_C(q_a = 0) = 0.$$

This is possible only when $p_a \leq 0$. Therefore, buyer C chooses $q^*_a = 1$ if $p_a \leq 0$ and chooses $q^*_a = 0$ otherwise.

In stage 1, seller A chooses price p_a to solve

$$\max_{p_a} \{p_a - \epsilon_a, 0\} \text{ s.t. } p_a \leq 0.$$

Because $p_a = 0$ yields negative payoff, seller A optimally chooses $p^*_a > \epsilon_a$. By Assumption 2(ii) ($\omega_c > \epsilon_b$), ensuring transaction costs are low enough, the subgame perfect equilibria are

$$\{p^*_a > \epsilon_a, (q^*_a = 0, q^*_{b1} = 1_c, q^*_{b0} = 0), (p^*_{b1} = \omega_c, p^*_{b0} > \epsilon_b)\}.$$

Along the equilibrium path, $\{p^*_a > \epsilon_a, (q^*_a = 0, q^*_{b0} = 0), p^*_{b0} > \epsilon_b\}$. Therefore, neither the composite good nor either component good is purchased in equilibrium. □

Proof of Proposition 2. *(No Refund/SAV_a/SAV_{ab} - No Tragedy)*

Assume $\gamma = 0$, $\omega_a > 0$, and $\omega_b \geq 0$. In stage 4 and then stage 3:

(i) if $q_a = 0$, then by Lemma 1, the Nash equilibrium of the second smallest subgame is $\{p^*_{b0} > \epsilon_b, q^*_{b0} = 0\}$ if $\omega_b = 0$ and $\{p^*_{b0} = \omega_b, q^*_{b0} = 1\}$ if $\omega_b > \epsilon_b$.

(ii) if $q_a = 1$, then in stage 4, by Assumption 1, buyer C never chooses $q_{b1} = 1_{a,b}$ for any $\omega_b \geq 0$. Buyer C chooses $q_{b1} = 1_c$ if and only if

$$\pi_C(q_{b1} = 1_c) \geq \pi_C(q_{b1} = 0) \Leftrightarrow \omega_c - p_a - p_b \geq \omega_a - p_a.$$

Therefore, by Assumption 3, the indifference rule, as long as $p_{b1} \leq \omega_c - \omega_a$, buyer C chooses $q_{b1} = 1_c$. In stage 3, seller B solves

$$\max_{p_{b1}} \{p_{b1} - \epsilon_b, 0\} \text{ s.t. } p_{b1} \leq \omega_c - \omega_a.$$

Therefore, seller B chooses $p^*_{b1} = \omega_c - \omega_a$.

Notice, assuming $\omega_b > \epsilon_b$, which holds by Assumption 2(i), then $\omega_b > 0$. Therefore, $\omega_c - \omega_a > \omega_b > \epsilon_b$, by Assumption 1. Otherwise, $\omega_b > 0$ holds by Assumption 2(ii) $(\omega_c - \omega_a > \epsilon_b)$.

In the third smallest subgame at stage 2, buyer C, knowing the optimal actions of seller B in stage 3, chooses $q_a \in \{0, 1\}$. By Assumption 3, the indifference rule, buyer C chooses to purchase component a as long as

$$\pi_C(q_a = 1) = \omega_c - p_a - p^*_{b1} = \omega_a - p_a \geq \pi_C(q_a = 0) = 0.$$

Therefore, buyer C chooses $q^*_a = 1$ if $p_a \leq \omega_a$ and chooses $q^*_a = 0$ otherwise.

In stage 1, seller A chooses p_a to solve

$$\max_{p_a} \{p_a - \epsilon_a, 0\} \text{ s.t. } p_a \leq \omega_a.$$

By Assumption 2(i) ($\omega_a > \epsilon_a$), then $p_a^* = \omega_a$. Therefore, assuming $\omega_c - \omega_a > \epsilon_b$ and $\omega_a > \epsilon_a$, the subgame perfect Nash equilibria are:

$$\{p_a^* = \omega_a, (q_a^* = 1, q_{b0}^* = 1, q_{b1}^* = 1_c), (p_{b0}^* = \omega_b, q_{b1}^* = \omega_c - \omega_a)\} \quad \text{if } SAV_{ab}$$
$$\text{and } \{p_a^* = \omega_a, (q_a^* = 1, q_{b0}^* = 0, q_{b1}^* = 1_c), (p_{b0}^* > 0, q_{b1}^* = \omega_c - \omega_a)\} \quad \text{if } SAV_a.$$

Notice, only off the equilibrium path do equilibria differ between the two games. For both games, along the equilibrium path,

$$\{p_a^* = \omega_a, (q_a^* = 1, q_{b1}^* = 1_c), q_{b1}^* = \omega_c - \omega_a\}.$$

Thus, the composite good is purchased in equilibrium. □

Proof of Proposition 3. *(No Refund/SAV_b)*

Assume $\gamma = 0$, $\omega_b > 0$, and $\omega_a = 0$. As explained in the text, a the game of no refund with stand alone value in only component b yields the same off-the-equilibrium path equilibria and equilibrium outcome of hold-up as the game of no refund with perfect complements. The only difference is the equilibria along the equilibrium path; since $\omega_b > 0$, then by Lemma 1, stand alone units of b are purchased in equilibrium. Therefore, the proof of Proposition 3 is the same as the proof of Proposition 1, except in the case where $q_a = 0$.

In stage 4 and then stage 3:

(i) if $q_a = 0$, then by Lemma 1, the Nash equilibrium of the second smallest subgame is $\{p_{b0}^* = \omega_b, q_{b0}^* = 1\}$.

(ii) if $q_a = 1$, the proof is identical to the proof of Proposition 1 for $q_a = 1$. Assuming $\omega_c > \epsilon_b$, the subgame perfect equilibria are

$$\{p_a^* > \epsilon_a, (q_a^* = 0, q_{b1}^* = 1_c, q_{b0}^* = 1), (p_{b1}^* = \omega_c, p_{b1}^* = \omega_b)\}.$$

Along the equilibrium path, $\{p_a^* > \epsilon_a, (q_a^* = 0, q_{b0}^* = 1), p_{b1}^* = \omega_b\}$. Therefore, the composite good is not purchased in equilibrium, but stand alone units of b are purchased in equilibrium. □

Proof of Lemma 2. *(B's pricing incentives under multiple quantities)*

Recall, $q_{b1} = 2_{c,b}$ means buyer C purchases two units of b, one to create composite c and one to use alone. Assume $\gamma = 0$, $\omega_b > 0$, and $\omega_a = 0$. Buyer C purchases two units of b as long as

$$\begin{aligned} \pi_C(q_{b1} = 2_{c,b}) &\geq \pi_C(q_{b1} = 1_c) \Leftrightarrow \\ (\omega_c - p_a - p_{b1}) + (\omega_b - p_{b1}) &\geq (\omega_c - p_a - p_{b1}) \Leftrightarrow \\ p_{b1} &\leq \omega_b \end{aligned} \tag{1.1}$$

and

$$\begin{aligned} \pi_C(q_{b1} = 2_{c,b}) &\geq \pi_C(q_{b1} = 0) \Leftrightarrow \\ (\omega_c - p_a - p_{b1}) + (\omega_b - p_{b1}) &\geq (-p_a) \Leftrightarrow \\ 2p_{b1} &\leq \omega_c + \omega_b. \end{aligned} \tag{1.2}$$

By Assumption 1, $\omega_c > \omega_b \Rightarrow \frac{\omega_c + \omega_b}{2} > \omega_b$. Therefore, condition (1.1), $p_{b1} \leq \omega_b$, is a necessary and sufficient condition to ensure buyer C purchases two units of b: one to create c and one to use alone.

Buyer C purchases one unit of b as long as

$$\begin{aligned} \pi_C(q_{b1} = 1_c) &> \pi_C(q_{b1} = 2_{c,b}) \Leftrightarrow \\ \omega_c - p_a - p_{b1} &> (\omega_c - p_a - p_{b1}) + (\omega_b - p_{b1}) \Leftrightarrow \\ \omega_b &< p_{b1} \end{aligned} \tag{1.3}$$

and

$$\begin{aligned} \pi_C(q_{b1} = 1_c) &\geq \pi_C(q_{b1} = 0) \Leftrightarrow \\ \omega_c - p_a - p_{b1} &\geq -p_a \Leftrightarrow \\ p_{b1} &\leq \omega_c. \end{aligned} \tag{1.4}$$

Therefore, by conditions (1.3) and (1.4), as long as $\omega_b < p_b \leq \omega_c$, buyer C purchases one unit of b to create composite c.

The highest price at which seller B sells two units is $p_{b1} = \omega_b$, yielding payoff $\pi_B(q_{b1} = 2_{c,b}) = 2\omega_b - \epsilon_b$. The highest price B sells one unit is $p_{b1} = \omega_c$, yielding payoff $\pi_B(q_{b1} = 1_c) = \omega_c - \epsilon_b$. Seller B chooses price $p_{b1} = \{\omega_c, \omega_b\}$ that yields the highest payoff. This price depends on the stand alone value of component b.

Recall, component b has *low stand alone value* if $\omega_c \geq 2\omega_b$ and *high stand alone value* if $\omega_c < 2\omega_b$. Therefore, by Assumption 2(i) ($\omega_c > \omega_b > \epsilon_b$), seller B prices to sell *two* units whenever b has *high stand alone value*:

$$\pi_B(q_{b1} = 2_{c,b}) > \pi_B(q_{b1} = 1_c) \quad \Leftrightarrow \quad 2\omega_b - \epsilon_b > \omega_c - \epsilon_b.$$[6]

[6]If transaction costs are marginal rather than fixed, then the result depends on high stand alone

Similarly, by Assumption 2(i) ($2\omega_b > \omega_b > \epsilon_b$), seller B prices to sell one unit whenever b has *low stand alone value*:

$$\pi_B(q_{b1} = 1_c) > \pi_B(q_{b1} = 2_{c,b}) \quad \Leftrightarrow \quad \omega_c - \epsilon_b > 2\omega_b - \epsilon_b.$$

$\square$

Proof of Proposition 4. *(No Refund/High SAV_b - No Tragedy, Low SAV_b - Tragedy)*

Assume $\gamma = 0$, $\omega_b > 0$, and $\omega_a = 0$.

part (i): In stage 4 and then stage 3:

(a) if $q_a = 0$, then by Lemma 1, the Nash equilibrium of the second smallest subgame is $\{p_{b0}^* = \omega_b, q_{b0}^* = 2\}$.

(b) if $q_a = 1$, then using the result and proof of Lemma 2, $p_{b1}^* = \omega_b$ and buyer C purchases $q_{b1} = 2_{c,b}$, using one b to create c and one to use alone.

In stage 2, buyer C faces profit of $\pi_C(q_a = 0) = 0$ and $\pi_C(q_a = 1) = \omega_c - \omega_b - p_a$. Therefore, as long as $p_a \leq \omega_c - \omega_b$, buyer C purchases $q_a = 1$.

In stage 1, seller A faces $\pi_A(p_a > \omega_c - \omega_b) = 0$ and $\pi_A(p_a \leq \omega_c - \omega_b) \leq \omega_c - \omega_b - \epsilon_a$. Therefore, by Assumption 2(ii) ($\omega_c - \omega_b > \epsilon_a$), to maximize profit, seller A chooses $p_a^* = \omega_c - \omega_b$.

The subgame perfect equilibrium, assuming $\omega_c > \epsilon_b$ and $\omega_c - \omega_b > \epsilon_a$ is:

$$\{p_a^* = \omega_c - \omega_b, (q_a^* = 1, q_{b1}^* = 2_{c,b}, q_{b0}^* = 1), (p_{b1}^* = \omega_b, p_{b0}^* = \omega_b)\}.$$

Along the equilibrium path, $\{p_a^* = \omega_c - \omega_b, (q_a^* = 1, q_{b1}^* = 2_{c,b}), (p_{b1}^* = \omega_b)\}$. Therefore, if b has *high stand alone value*, the composite good is purchased in equilibrium.

value in b *along with* the marginal benefit to supplying the second unit greater than the marginal cost to supplying the second unit. I leave the result in terms of fixed costs to keep focus on stand alone value.

part (ii): (a) if $q_a = 0$, then by Lemma 1, the Nash equilibrium of the second smallest subgame is $\{p^*_{b0} = \omega_b, q^*_{b0} = 1\}$.

(b) if $q_a = 1$, then in stage 4, buyer C purchases $q_{b1} = 1_c$ as long as $\omega_b \leq p_{b1} \leq \omega_c$. Buyer C purchases $q_{b1} = 2_{c,b}$ as long as $p_{b1} \leq \omega_b$ and $2p_{b1} \leq \omega_c + \omega_b$. Since $\frac{\omega_c+\omega_b}{2} > \omega_b$, buyer C purchases $q_{b1} = 2_{c,b}$ as long as $p_{b1} \leq \omega_b$.

Lemma 2 shows that when component b has *low stand alone value*, seller B prices at the monopoly price for composite good c. Therefore, by Assumption 2(i) ($\omega_c > \omega_b > \epsilon_b$), seller B chooses $p^*_b = \omega_c$.

The remainder of the proof follows the proof of Proposition 1 (and hence, the proof of Proposition 3).

In stage 2, buyer C faces profit $\pi_C(q_a = 0) = 0$ and $\pi_C(q_a = 1) = -p_a$. As long as seller A prices at $p_a \leq 0$, then buyer C purchases component a. However, in stage 1, by the assumption of transaction costs, at price at $p_a \leq 0$, seller A's surplus is negative. Therefore, seller A chooses $p^*_a > 0$ to maximize surplus.

Assuming b has low stand alone value, the subgame perfect equilibria, assuming $2\omega_b > \epsilon_b$, are:

$$\{p^*_a > 0, (q^*_a = 0, q^*_{b1} = 1_c, q^*_{b0} = 2), (p^*_{b1} = \omega_c, p^*_{b0} = \omega_b)\}.$$

Along the equilibrium path, component a is not purchased. Therefore, if b has low stand alone value, the composite good is not purchased in equilibrium, but stand alone units of b (two, in this case) are purchased in equilibrium. □

Proof of Proposition 5. *(No Refund/SAV$_b$: Price Discrimination - Tragedy)*

Assume $\gamma = 0$, $\omega_b > 0$, and $\omega_a = 0$. Let p_{b1c} be the price B chooses, corresponding to the purchase of b to use to form c. Let p_{b1b} be the price B chooses, corresponding to the purchase of b to use alone. In stage 4 and then stage 3:

(i) if $q_a = 0$, then by a similar result to Lemma 1 (but extending $q_{b0} \in \{0, 1, 2\}$), the Nash equilibrium of the second smallest subgame is $\{p^*_{b0} = \omega_b, q^*_{b0} = 2\}$.

(ii) if $q_a = 1$, in stage 4, by Assumption 3, buyer C purchases $q_{b1} = 1_c$ as long as $p_{b1c} \leq \omega_c$ and $p_{b1b} \geq \omega_b$. Buyer C purchases $q_{b1} = 2_{c,b}$ as long as $p_{b1b} \leq \omega_b$ and $p_{b1c} + p_{b1b} \leq \omega_c + \omega_b$.

Therefore, by Assumption 1 and Assumption 2(i) ($\omega_c > \omega_b > \epsilon_b$) in stage 3, seller B chooses $p^*_{b11} = \omega_c$ for the first unit of b, since seller B knows buyer C prefers to use this unit with the one-unit of a purchased in stage 2. For the second unit of b purchased by buyer C, seller B can price up to buyer C's reservation value of using b alone. Therefore, by Assumption 2(i) ($\omega_b > \epsilon_b$), $p^*_{b12} = \omega_b$.

Based on these optimal pricing decisions by seller B, buyer C's payoffs for any number of units of b purchased is $-p_a$. Therefore, in stage 2, by Assumption 3, the indifference rule, buyer C purchases $q_a = 1$ only if $p_a \leq 0$. However, $p_a \leq 0$ is suboptimal for seller A given transaction costs are positive. Therefore, A chooses $p^*_a > 0$ to maximize surplus.

The subgame perfect equilibria, assuming $\omega_b > \epsilon_b$, are:

$$\{p^*_a > 0, (q^*_a = 0, q^*_{b1} = 1_c, q^*_{b0} = 2), (p^*_{b11} = \omega_c, p^*_{b12} = \omega_b, p^*_{b0} = \omega_b)\} .$$

Along the equilibrium path, component a is not purchased. Therefore, in equilibrium, the composite good is not purchased, but stand alone units of b (two in this case) are purchased. □

Proofs of Proposition 6 and Proposition 7. *(Partial and Full Refunds)*

Each proof of full refund builds from the proofs of partial refund. Therefore, I complete the proofs for partial and full refunds in the following order: (I.a) and (I.b): perfect complements; (II.a) and (II.b): SAV_b; (III.a) and (III.b): SAV_a; and (IV.a) and (IV.b): SAV_{ab}.

(I.a) Perfect complements and partial refund: Assume $\omega_b = 0$, $\omega_a = 0$ and $0 < \gamma < 1$. In stage 4 and then stage 3:

(i) if $q_a = 0$, then by Lemma 1, the Nash equilibria of the second smallest subgame is $\{p^*_{b0} > \epsilon_b, q^*_{b0} = 0\}$.

(ii) if $q_a = 1$, then in stage 4, buyer C never chooses $q_{b1} = 0$ or $q_{b1} = 1_{a,b}$. By Assumption 3, the indifference rule, buyer C chooses $q_{b1} = 1_c$ as long as

$$\omega_c - \gamma p_a \geq p_{b1} \quad \text{and} \quad \omega_c - \gamma p_a > 0.$$

Buyer C chooses $q_{b1} = 1_{\not{c},b}$ as long as

$$p_{b1} \leq 0 \quad \text{and} \quad \omega_c - \gamma p_a < 0.$$

In stage 3, buyer B maximizes payoffs given the optimal behavior of buyer C in stage 4. Seller B faces

$$\pi_B(q_{b1} = 1_c) = p_{b1} - \epsilon_b = \omega_c - \gamma p_a - \epsilon_b,$$

$$\pi_B(q_{b1} = 0_{\not{c}}) = 0, \quad \text{and}$$

$$\pi_B(q_{b1} = 1_{\not{c},b}) = p_{b1} = \epsilon_b = -\epsilon_b.$$

Therefore, seller B maximizes payoffs choosing

$$p_{b1}^{*} = \begin{cases} \omega_c - \gamma p_a & \text{if } \omega_c - \gamma p_a - \epsilon_b > 0 \\ 0 & \text{if } \omega_c - \gamma p_a - \epsilon_b \leq 0 \end{cases}$$

In stage 2, buyer C faces the following payoffs:

$$\text{if } q_a = 0\text{: } \pi_C = 0$$

and

$$\text{if } q_a = 1\text{: } \begin{cases} \pi_C(q_{b1} = 1_c) = \omega_c - p_a - p_{b1}^{*} = \omega_c - p_a - \omega_c + \gamma p_a = p_a(\gamma - 1) \\ \qquad \text{if } \omega_c - \gamma p_a - \epsilon_b > 0 \quad \text{ and} \\ \pi_C(q_a = 1, q_{b1} = 0_{\not{c}}) = p_a(\gamma - 1) \\ \qquad \text{if } \omega_c - \gamma p_a - \epsilon_b \leq 0. \end{cases}$$

Therefore, as long as $p_a(\gamma - 1) \geq 0$, then buyer C chooses $q_a = 1$. Under a partial refund, $p_a(\gamma - 1) \geq 0$ only for $p_a \leq 0$.

In stage 1, seller A solves

$$\max_{p_a} \{p_a - \epsilon_a, 0\} \quad \text{s.t} \quad p_a \leq 0.$$

Therefore, since there exists no positive price for which seller A achieves positive payoffs, seller A chooses $p_a^{*} > \epsilon_a$.

The subgame perfect equilibria for the case of partial refund and perfect complementarities are:

$$\text{SPNE} = \begin{cases} p_a^* > \epsilon_a & \\ (q_a^* = 0, q_{b0}^* = 0, q_{b1}^* = 1_c) & \text{if } p_{b1}^* \leq \omega_c - \gamma p_a^* \text{ and } \gamma p_a^* < \omega_c. \\ (q_a^* = 0, q_{b0}^* = 0, q_{b1}^* = 0) & \text{if } p_{b1}^* = 0 \text{ and } \omega_c < \gamma p_a^* \\ (p_{b0}^* > \epsilon_b, p_{b1}^* = \omega_c - \gamma p_a^*) & \text{if } \gamma p_a^* < \omega_c - \epsilon_b \\ (p_{b0}^* > \epsilon_b, p_{b1}^* = 0) & \text{if } \gamma p_a^* > \omega_c - \epsilon_b \end{cases}$$

Along the equilibrium path,

$$\{p_a^* > \epsilon_a, q_{b0}^* = 0, p_{b0}^* > \epsilon_b\}.$$

In equilibrium, nether the composite nor either component is purchased. Equilibrium payoffs are $\pi_A^* = 0, \pi_C^* = 0, \pi_B^* = 0$.

(I.b) Perfect complements and full refund: Now assume $\gamma = 1$. The proof follows that from (I.a) up to stage 2. At stage 2, recall as long as $p_a(\gamma - 1) \geq 0$ then buyer C chooses $q_a = 1$. Under full refund, this condition holds for all $p_a \geq 0$. In stage 1, seller A chooses $p_a \geq 0$; however, seller A must price to ensure seller B prices to sell b so buyer C chooses $q_{b1} = 1_c$ and not $q_{b1} = 1_{\not{a},b}$. If seller A prices higher than $\omega_c - \epsilon_b$ then seller B's profits are negative. Therefore, per seller B's optimal stage 3 behavior, seller A chooses $p_a < \omega_c - \epsilon_b$. Let $\delta > 0$ be the added surplus seller A leaves on the table for seller B to optimally price so C chooses $q_a = 1_c$ and to ensure seller B is not indifferent. Then seller A chooses $p_a^* = \omega_c - \delta - \epsilon_b$ as long as $\omega_c - \delta - \epsilon_b - \epsilon_a > 0$.

Therefore, the subgame perfect Nash equilibrium under a full refund and perfect complements, assuming $\delta > 0$ and $\omega_c - \delta - \epsilon_b - \epsilon_a > 0$ is:

$$\{p_a^* = \omega_c - \delta - \epsilon_b, (q_a^* = 1, q_{b0}^* = 0, q_{b1}^* = 1_c), (p_{b0}^* > \epsilon_b, p_{b1}^* = \epsilon_b + \delta)\}.$$

Therefore, in equilibrium, buyer C purchases the composite good. Equilibrium payoffs are:

$$\pi_A^* = \omega_c - \epsilon_b - \delta - \epsilon_a, \quad \pi_C^* = 0, \quad \pi_B^* = \delta.$$

Therefore, in equilibrium, when no component has stand alone value, for any refund less than full (zero or partial), buyer C does not purchase the composite good, and for a full refund, buyer C purchases the composite good.

(II.a) SAV_b **and partial refund:** Assume $\omega_b > 0$, $\omega_a = 0$, and $0 < \gamma < 1$. In stage 4 and then stage 3:

(i) if $q_a = 0$, then by Lemma 1, the Nash equilibria of the second smallest subgame is $\{p_{b0}^* = \omega_b, q_{b0}^* = 1\}$.

(ii) if $q_a = 1$, then by Assumption 3, the indifference rule, in stage 4, buyer C purchases $q_b = 1$ and keeps component a as long as $p_b \leq \omega_c - \gamma p_a$ and $\omega_c - \omega_b > \gamma p_a$. Buyer C purchases $q_b = 1$ and returns component a whenever $p_b \leq \omega_b$ and $\omega_c - \omega_b < \gamma p_a$.

In stage 3, seller B chooses p_b to maximize surplus. Let p_{bb_1} and p_{bb_2} be two prices seller B chooses to offer, determined in the following way. The highest price seller B can offer when $\omega_c - \omega_b > \gamma p_a$ is $p_{bb_1} = \omega_c - \gamma p_a$. The highest price seller B can offer when $\omega_c - \omega_b < \gamma p_a$ is $p_{bb_2} = \omega_b$. Seller B chooses $p_{bb_1} = \omega_c - \gamma p_a$ whenever $\omega_c - \omega_b > \gamma p_a$, yielding $\pi_B(q_b^{keep_a}) = \omega_c - \gamma p_a - \epsilon_b$. Seller B chooses $p_{bb_2} = \omega_b$ whenever $\omega_c - \omega_b < \gamma p_a$, yielding $\pi_B(q_b^{ret_a}) = \omega_b - \epsilon_b$. Seller B's optimal price $p_{b1}^* \in \{p_{bb_1}, p_{bb_2}\}$ depends on the price chosen by seller A.

Under a partial refund assumption, in stage 2, buyer C chooses $q_a = \{0, 1\}$ to maximize surplus. Buyer C faces $\pi_C(q_a = 1) = \gamma p_a - p_a$ and $\pi_C(q_a = 0) = 0$. By the indifference rule, as long as $\gamma p_a - p_a = p_a(\gamma - 1) \geq 0$, then buyer C chooses $q_a = 1$. Clearly, unless $\gamma = 1$, this inequality does not hold. Therefore, for any $\gamma \in [0, 1)$, when component b has stand alone value, regardless of the price chosen by seller A, buyer C never chooses to purchase component a, resulting in hold-up. The subgame perfect Nash equilibria in the case of positive SAV_b and partial refund are:

$$\text{SPNE} = \begin{cases} p_a^* > \epsilon_a & \\ (q_a^* = 0, q_{b0}^* = 1, q_{b1}^* = 1_c) & \text{if } p_{b1}^* \leq \omega_c - \gamma p_a^* \text{ and } \gamma p_a^* < \omega_c - \omega_b. \\ (q_a^* = 0, q_{b0}^* = 0, q_{b1}^* = 1_{\not{c}b}) & \text{if } p_{b1}^* \leq \omega_b \text{ and} \omega_c - \omega_b < \gamma p_a^*) \\ (p_{b0}^* = \omega_b, p_{b1}^* = \omega_c - \gamma p_a^*) & \text{if } \gamma p_a^* < \omega_c - \omega_b \\ (p_{b0}^* = \omega_b, p_{b1}^* = \omega_b) & \text{if } \gamma p_a^* > \omega_c - \omega_b \end{cases}$$

In equilibrium, composite good c is not purchased, but stand alone units of b are purchased. Equilibrium payoffs are

$$\pi_A^* = 0, \quad \pi_C^* = 0, \quad \pi_B^* = \omega_b - \epsilon_b.$$

(II.b) SAV_b **and full refund:** Now assume $\gamma = 1$. Under a full refund assumption, if seller A offers a full refund, then in stage 2, when buyer C chooses $q_a = \{0, 1\}$, by the indifference rule, buyer C chooses $q_a = 1$ for any $p_a \geq 0$. In order for seller A to maximize profits, seller A chooses p_a in such a way that seller B chooses p_{b_1} (buyer C keeps a) rather than p_{b_2}. Therefore, seller A chooses $p_a^* < \omega_c - \omega_b$ to ensure seller B's payoff is highest when buyer C keeps component a. Assuming p_a^* is δ smaller than $p_a^* = \omega_c - \omega_b$, then seller B chooses $p_{b1} = \omega_c - p_a^* = \omega_b + \delta$. Seller A must leave just more than ω_b on the table to give seller B incentive to price to sell $q_{b1} = 1_c$ rather than $q_{b1} = 1_{\not{a},b}$.

In equilibrium, assuming $\omega_c - \omega_b - \epsilon_a > \delta > 0$,

$$\{p_a^* = \omega_c - \omega_b - \delta, (q_a^* = 1, q_{b0}^* = 1, q_{b1}^* = 1_c), (p_{b0}^* = \omega_b, p_{b1}^* = \omega_b + \delta)\}$$

Composite good c is purchased in equilibrium. Equilibrium payoffs are:

$$\pi_A^* = \omega_c - \epsilon_b - \delta - \epsilon_a, \quad \pi_C^* = 0, \quad \pi_B^* = \delta.$$

Therefore, in equilibrium, when only component b has stand alone value, for any refund less than full (zero or partial), buyer C does not purchase the composite good, and for a full refund, buyer C purchases the composite good.

(III.a) SAV_a **and partial refund:** Assume $\omega_a > 0$, $\omega_b = 0$, and $0 < \gamma < 1$. In stage 4 and then stage 3:

(i) if $q_a = 0$, then by Lemma 1, the Nash equilibria of the second smallest subgame is $\{p_{b0}^* > \epsilon_b, q_{b0}^* = 0\}$.

(ii) if $q_a = 1$, in stage 4, buyer C chooses $q_{b1} = 1_c$ as long as $p_{b1} \leq \omega_c - \max\{\omega_a, \gamma p_a\}$ and $\gamma p_a < \omega_c$. In stage 3, seller B chooses $p_{b1} > \epsilon_b$ to maximize profit. To sell one unit of b, $p_{b1} \in \{0, \omega_c - \max\{\omega_a, \gamma p_a\}\}$. Since $\pi_B(p_{b1} = 0) = -\epsilon_b$, seller B would always choose to sell zero units of b than price at zero. Therefore, as long as $\omega_c - \max\{\omega_a, \gamma p_a\} - \epsilon_b > 0$ then seller B chooses $p_b^* = \omega_c - \max\{\omega_a, \gamma p_a\}$ so buyer C purchases b to form c.

In stage 2, by the indifference rule, buyer C chooses $q_a = 1$ as long as $p_{b1} = \omega_c - \max\{\omega_a, \gamma p_a\}$ and $\pi_C(q_a = 1, q_{b1} = 1_c) = \max\{\omega_a, \gamma p_a\} - p_a \geq 0$.

In stage 1, seller A chooses p_a to maximize payoffs. Seller A prices to sell a as long as $p_a^* > \epsilon_a$. If seller A chooses $p_a = \omega_a$ then buyer C purchases and keeps a. However, since buyer C values composite c more than component a, seller A can price higher than ω_a and still sell a. Seller A chooses p_a so buyer C chooses $q_a = 1$ and seller B chooses p_{b1} so that $q_{b1} = 1_c$. This means p_a^* must satisfy (1) $\max\{\omega_a, \gamma p_a^*\} \geq p_a^*$; (2) $\omega_c - \max\{\omega_a, \gamma p_a^*\} - \epsilon_b > 0$; and (3) $\gamma p_a \leq \omega_c$.

In the case of partial refund, seller A must choose p_a so $\omega_a = \max\{\omega_a, \gamma p_a^*\}$; otherwise, there exists no p_a such that condition (1) is satisfied. The only p_a for which condition (1) holds is $p_a = \omega_a$. Assuming ϵ_b is small enough, for $p_a = \omega_a$, condition (2) is satisfied. Condition (3) is satisfied since $\omega_a < \omega_c$. Therefore, assuming ϵ_a is small enough, A maximizes payoffs by choosing $p_a^* = \omega_a$.

Under partial refund, by Assumption 2(i) and (ii), ($\omega_a > \epsilon_a$ and $\omega_c - \omega_a > \epsilon_b$), the subgame perfect equilibrium under a partial refund when only component a has stand alone value is

$$\{p_a^* = \omega_a, (q_a^* = 1, q_{b0}^* = 0, q_{b1}^* = 1_c), (p_{b0}^* > 0, p_{b1}^* = \omega_c - \omega_a)\}.$$

The composite good is purchased in equilibrium under partial refund, and the equilibrium payoffs are

$$\pi_A^* = \omega_a - \epsilon_a,\ \pi_C^* = 0,\ \pi_B^* = \omega_c - \omega_a - \epsilon_b.$$

(III.b) SAV_a **and full refund:** Now assume $\gamma = 1$. In the case of full refund, condition (1), above, is satisfied for any p_a. Therefore, it is possible for seller A to price higher than ω_a, but condition (2), and hence condition (3), are satisfied if and only if $p_a < \omega_c - \epsilon_b$. By Assumption 2(ii) where $\omega_c - \omega_a > \epsilon_b$, then for $\delta < \omega_c - \epsilon_b$, seller A maximizes payoff by choosing $p_a^* = \omega_c - \epsilon_b - \delta$.

Now, $\omega_c - \epsilon_b - \delta > \epsilon_a$ implies $\delta < \omega_c - \epsilon_b$. Under full refund, by Assumption 2(ii) ($\omega_c - \epsilon_b > \omega_a$) and assuming $\omega_c - \epsilon_b - \epsilon_a > \delta > 0$, the subgame perfect equilibrium is

$$\{p_a^* = \omega_c - \epsilon_b - \delta, (q_a^* = 1, q_{b0}^* = 0, q_{b1}^* = 1_c), (p_{b0}^* > \epsilon_b, p_{b1}^* = \epsilon_b + \delta)\}.$$

The composite good is purchased in equilibrium, and the equilibrium payoffs are

$$\pi_A^* = \omega_c - \epsilon_b - \delta - \epsilon_a,\ \pi_C^* = 0,\ \pi_B^* = \delta.$$

Therefore, in equilibrium, when only component a has stand alone value, for any refund (zero, partial or full), buyer C purchases the composite good.

(IV.a) SAV_{ab} and partial refund: Assume $\omega_a > 0$, $\omega_b > 0$, and $0 < \gamma < 1$. When both components have stand alone value, the proof for a partial refund is similar to that when only component a has stand alone value. The stand alone value in a is enough to overcome hold-up by seller B, and the partial refund only affects equilibrium payoffs. The subgame perfect Nash equilibrium along the equilibrium path is the same as the case of SAV_a:

$$\{p_a^* = \omega_a, (q_a^* = 1, q_{b0}^* = 1, q_{b1}^* = 1_c), (p_{b0}^* = \omega_b, p_{b1}^* = \omega_c - \omega_a)\}$$

Composite c is purchased in equilibrium, and equilibrium payoffs are

$$\pi_A^* = \omega_a - \epsilon_a, \quad \pi_C^* = 0, \quad \pi_B^* = \omega_c - \omega_a - \epsilon_b.$$

(IV.b) SAV_{ab} and full refund: Now assume $\gamma = 1$. When both components have stand alone value and a full refund is offered, the proof is similar to that of stand alone value in only component b. Seller A must account for the fact that stand alone value in b provides leverage to seller B that prevents seller A from taking all of buyer C's surplus. Seller A must leave enough surplus on the table for seller B to have incentive to price to sell b so buyer C does not return a ($q_{b1} = 1_c$) rather than C purchase b to use alone and return a (i.e does not choose $q_{b1} = 1_{\not{a},b}$). In equilibrium, assuming $\omega_c - \omega_b - \epsilon_a > \delta > 0$,

$$\{p_a^* = \omega_c - \omega_b - \delta, (q_a^* = 1, q_{b0}^* = 1, q_{b1}^* = 1c), (p_{b0}^* = \omega_b, p_{b1}^* = \omega_b + \delta)\}$$

Equilibrium payoffs are:

$$\pi_A^* = \omega_c - \epsilon_b - \delta - \epsilon_a, \quad \pi_C^* = 0, \quad \pi_B^* = \delta$$

and the composite good is purchased in equilibrium. □

Therefore, in equilibrium, when both components a and b have stand alone value, for any refund (zero, partial or full), buyer C purchases the composite good.

Proof of Corollary 1. Seller A chooses a full refund if the equilibrium payoff to seller A under a zero or partial refund is lower than the equilibrium payoff to seller A under a full refund.

(i) As long as $\omega_c - \omega_a > \delta$ then

$$\pi_A(SAV_a, \gamma = 1) = \omega_c - \delta - \epsilon_a \quad > \quad \pi_A(SAV_a, \gamma < 1) = \omega_a - \epsilon_a.$$

Since $\delta > \epsilon_b$ when component a has stand alone value, then either seller B's transaction costs cannot be too high or component a's stand alone value cannot be too high.

(ii) As long as $\omega_c - (\omega_a + \omega_b) > \delta$ then

$$\pi_A(SAV_{ab}, \gamma = 1) = \omega_c - \omega_b - \delta - \epsilon_a \quad > \quad \pi_A(SAV_a, \gamma < 1) = \omega_a - \epsilon_a.$$

Therefore, either the net value of the composite must be greater than the amount of surplus seller A must leave on the table for seller B, or the values of components a and b cannot be too high.

For (i) and (ii), if the gains from offering a full refund are greater than the surplus foregone in order to ensure B prices to sell b with a, then seller A prefers a full refund. □

Proof of Corollary 2. The equilibrium outcome of both a game of SAV_a and game of SAV_{ab} yields equilibrium surplus

$$\pi_A^* = \omega_a - \epsilon_a \quad \text{and} \quad \pi_B^* = \omega_c - \omega_a - \epsilon_b.$$

By definition of *high (low) stand alone value* in component a, $\omega_c < (>) 2\omega_a$. Assuming transaction costs are equivalent, $\epsilon_a = \epsilon_b = \epsilon$, then

$$\begin{aligned} \omega_c &< (>) \;\; 2\omega_a \quad \Leftrightarrow \\ \omega_c - \omega_a &< (>) \;\; \omega_a \quad \Leftrightarrow \\ p_{b1}^* &< (>) \;\; p_a^* \quad \Leftrightarrow \\ \pi_B^* + \epsilon_b &< (>) \;\; \pi_A^* + \epsilon_a \quad \Leftrightarrow \\ \pi_B^* &< (>) \;\; \pi_A^* \end{aligned}$$

which proves the claim. □

Proof of Corollary 3. If seller A offers a full refund and at least component b has stand alone value, then seller B realizes equilibrium payoffs

$$\pi_B^* = \omega_b + \delta - \epsilon_b$$

while seller A realizes equilibrium payoffs

$$\pi_A^* = \omega_c - \omega_b - \delta - \epsilon_a.$$

Therefore,

$$\begin{aligned} \pi_B^* &> \pi_A^* && \Leftrightarrow \\ \omega_b + \delta - \epsilon_b &> \omega_c - \omega_b - \delta - \epsilon_a && \Leftrightarrow \\ (\delta + \epsilon_a) + (\delta - \epsilon_b) &> \omega_c - 2\omega_b. \end{aligned}$$

Assuming $\delta > \epsilon_b$, the first term above is positive. Assuming high stand alone value in b, the second term is negative. Therefore, the result holds. □

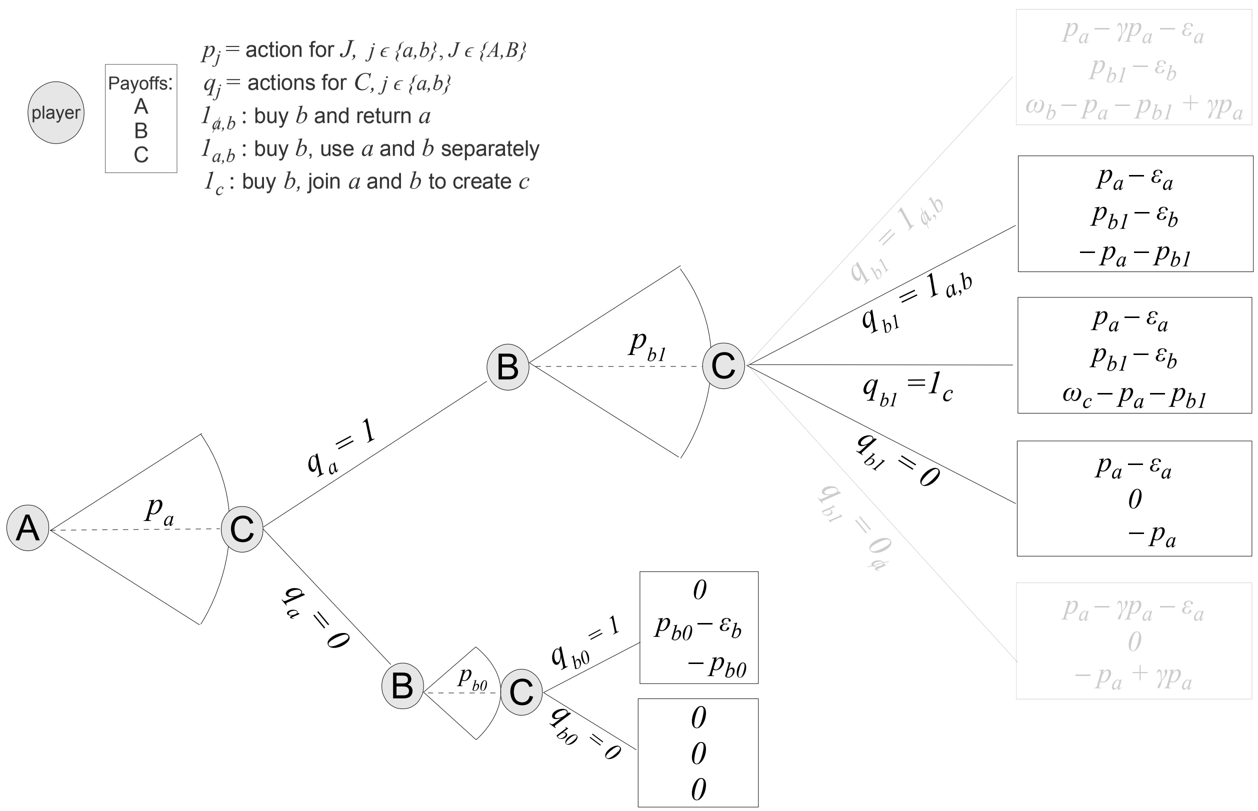

Figure 1.6. Game of no refund with perfect complements

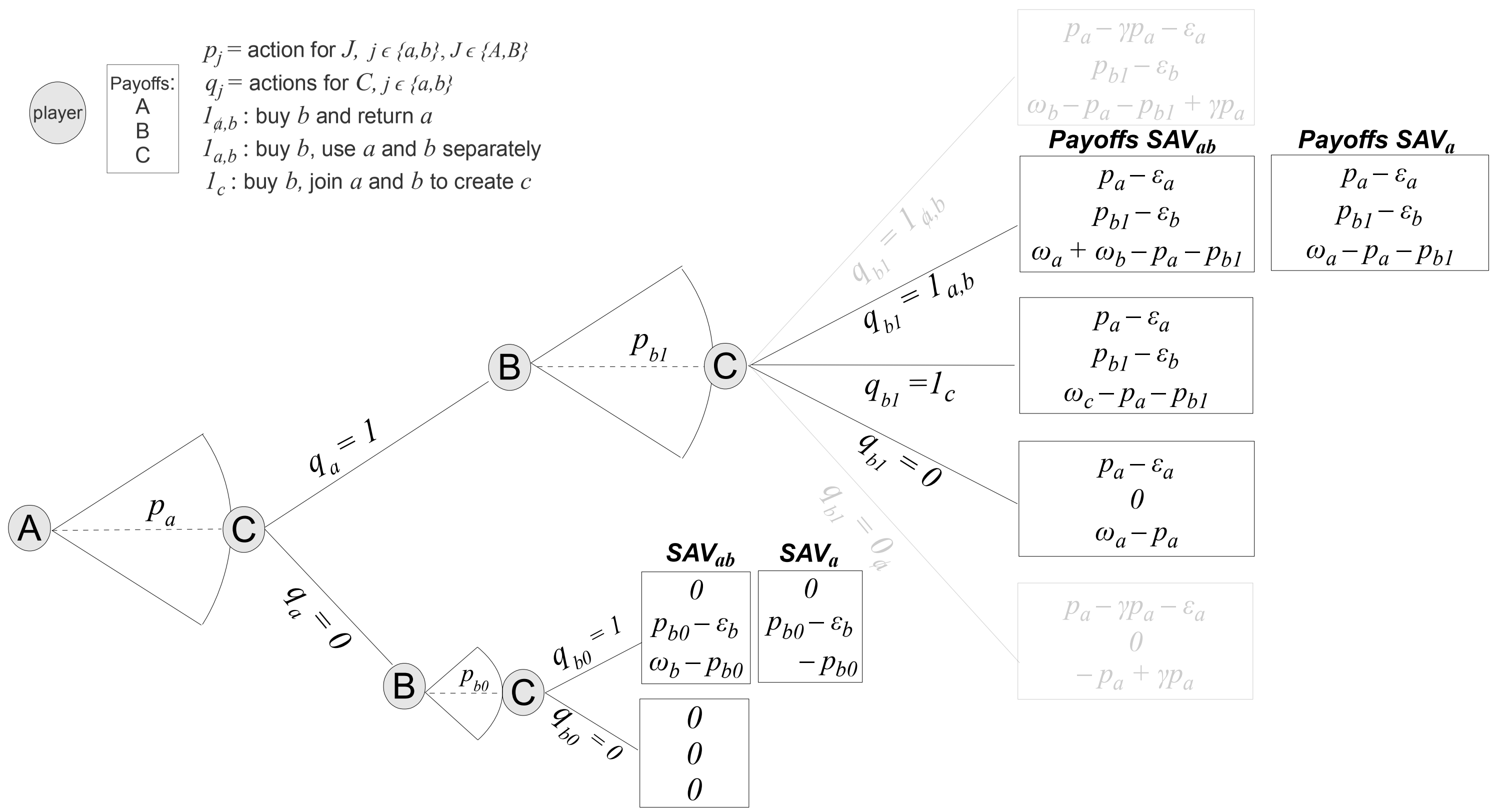

FIGURE 1.7. Game of no refund with SAV_{ab} and SAV_a

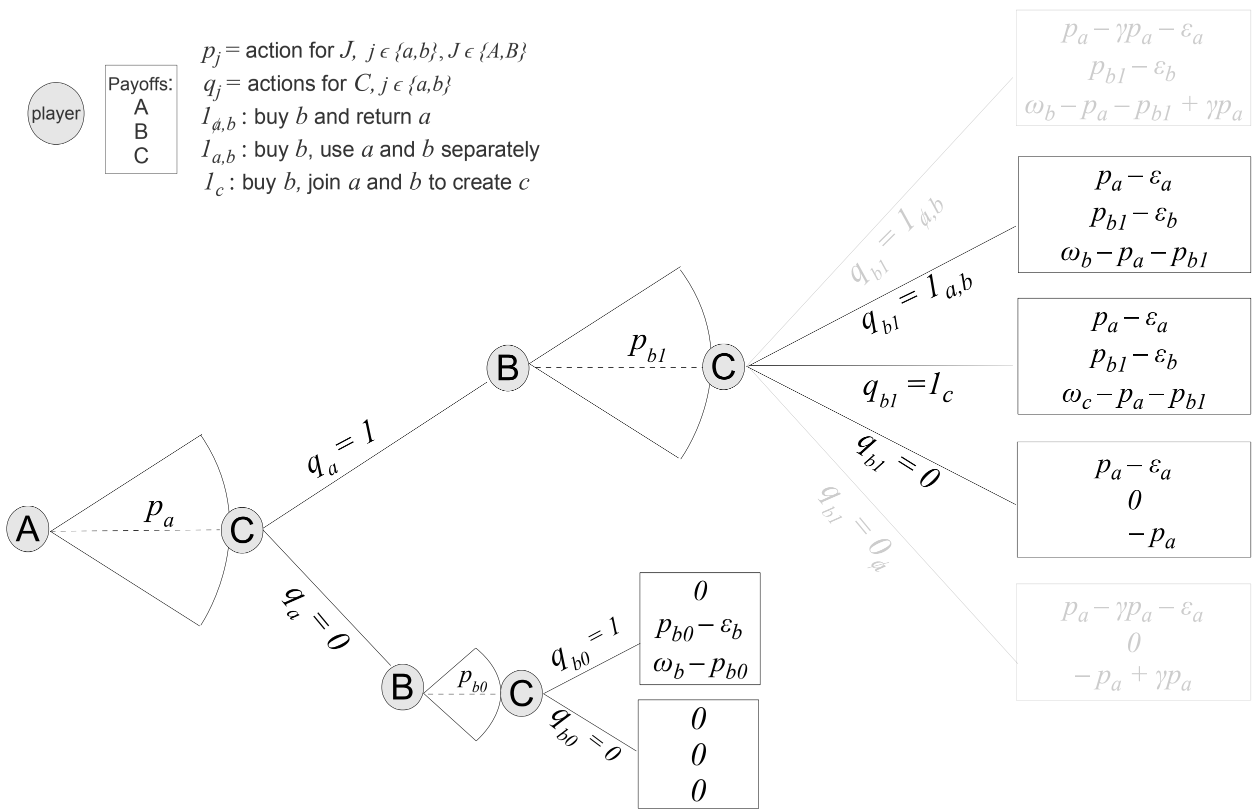

FIGURE 1.8. Game of no refund and SAV_b

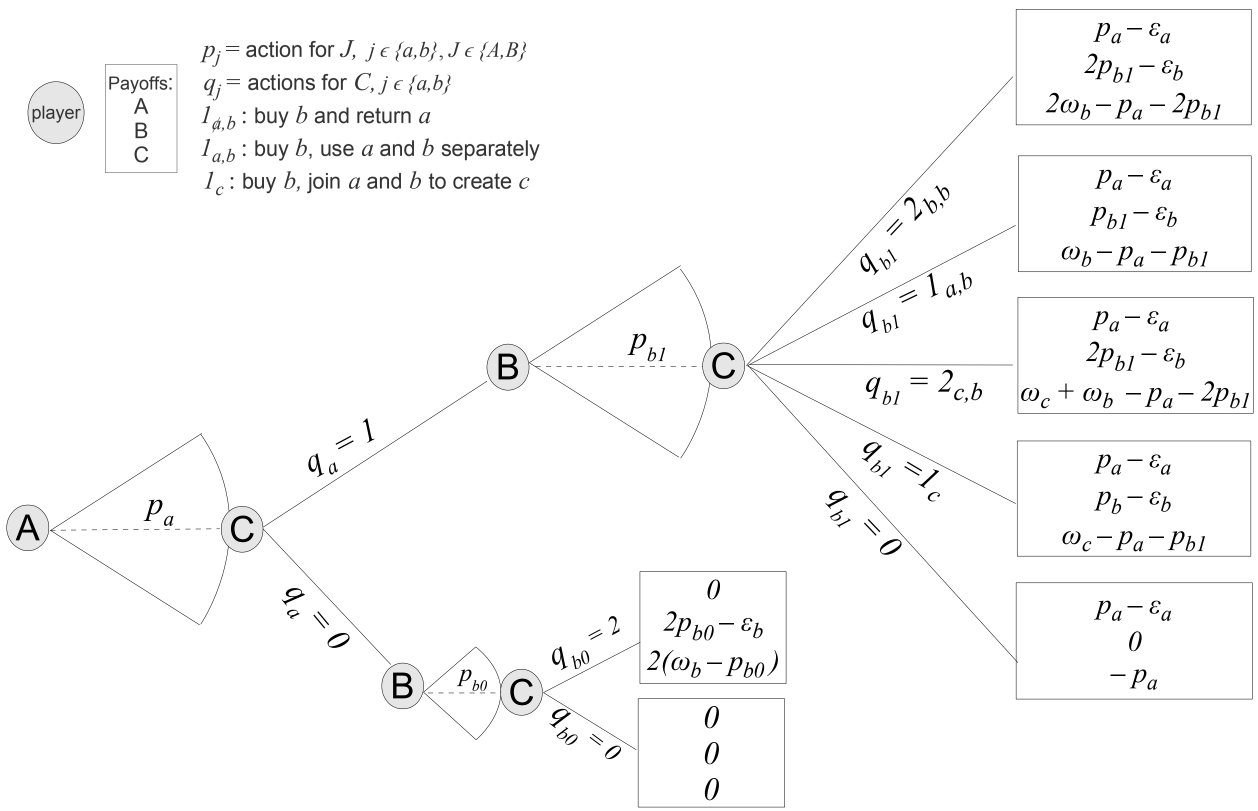

FIGURE 1.9. Game of no refund and SAV_b with unconstrained capacity in b and no price discrimination

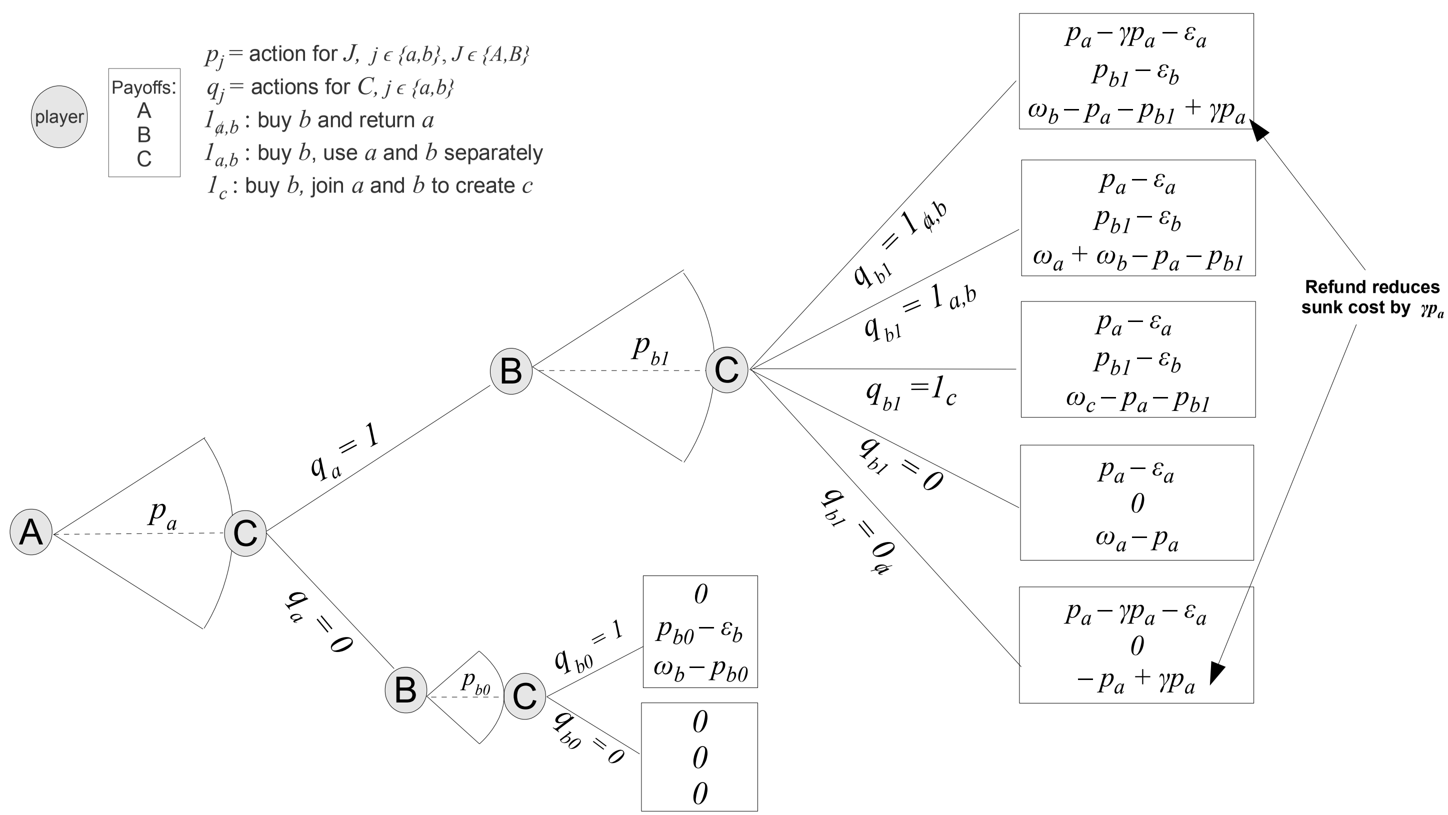

FIGURE 1.10. Game of refund $\gamma \in [0, 1]$ with nonnegative SAV_a and SAV_b

Chapter 2
Bayesian Analysis of a Single-Firm Event Study

> "In a general sense, the most attractive feature of the Bayesian approach is that it encourages one to direct energy toward "getting the model right" and away from modeling compromises required to utilize standard inferential methodology." -Gelfand and Sfiridis (1996)

1. Introduction

Event studies, commonly used in securities litigation or antitrust investigations, examine the effects of one or several events on a firm's stock return. In securities litigation, single-firm event studies analyze the impact on stock prices of alleged fraudulent behavior, such as corrective or misleading disclosure statements. In antitrust investigations, an event study can reveal the financial markets' expectations concerning the impact of a merger on firms' market valuations. An event study seeks to determine what would have happened to a firm's stock price "but for" the alleged fraud or potential merger.[1]

The standard event study methodology uses a frequentist approach to inference based on an assumption of normally distributed stock returns. Such a method is advocated, if not mandated, by courts in securities litigation to establish materiality and loss causation, as well as to determine the amount of loss.[2] In *Basic v. Levinson*

[1] We examine event studies concerning stock prices. However, bond prices also could be analyzed. Furthermore, while we focus primarily on securities litigation, our analysis and results apply to antitrust and merger event studies. See Cichello and Lamdin (2006) for an overview of event studies applied to antitrust.

[2] For example, *In re Omnicom Group, Inc. Securities Litigation (2010)* turned on the improper

(1988), the Supreme Court ruled that if a plaintiff establishes that the market for the alleged firm's security sufficiently provides all material information such that the security's price is quickly reflected in any public statements, then the reliance element of SEC Rule 10b-5 is met. Proving that the alleged misleading or omitted information would have affected the market's valuation of the security is a question of fact, so event studies are used to establish reliance.

Acceptance of a single-firm event study in securities litigation stems from court rulings finding that it satisfies what has become known as the *Daubert* standard. In *Daubert v. Merrell Dow Pharmaceuticals, Inc. (1993)*, the Supreme Court ruled that when expert testimony uses scientific evidence, the scientific technique or theories used to exhibit evidence must (i) be subjected to peer review; (ii) be found valid; (iii) be generally accepted; and (iv) yield a statistical error rate low enough to ensure reliable results. *Kumho Tire Co. v. Carmichael (1999)* extended *Daubert* to non-scientific expert testimony, thereby treating economic, accounting and financial experts to the same standards. Until recent scrutiny, standard inference for single-firm event studies has been deemed to satisfy the *Daubert* standard.

Classical single-firm event studies use standard inference to analyze a firm's excess stock return from an event day (often called the event effect). With only one event date, the estimated event effect cannot be viewed as the average of a large number of random variables. Therefore, to be valid, standard inference requires the assumption that excess returns, also known as abnormal returns, are normally distributed Gelbach, Helland, and Klick (2010)). Unless excess returns are exactly normally dis-

use of the event study methodology, and in *Oscar Private Equity Investments v. Allegiance Telecom, Inc. (2007)* the District Court found the expert's event study "untenable." Additionally, *In re Williams Securities Litigation-WCG Subclass (2009)*, the Tenth Circuit Court of Appeals affirmed the lower court's decision to exclude the plaintiff's expert testimony for failure to perform an event study, and in *In re Imperial Credit Industries, Inc. (2008)*, the court dismissed the plaintiffs' expert testimony, which did not include an event study, granting summary judgment to the defendants.

tributed, then standard inference is statistically invalid, in which case the *Daubert* standard is not satisfied. Since establishing loss under standard methods of inference may invoke invalid assumptions, alternative methods not only are sought after by scholars, but also are important in practice.[3]

One alternative to the standard frequentist approach is a Bayesian analysis of an event study. Few have analyzed the event study methodology through a Bayesian perspective, and none consider a Bayesian analysis of a single-firm event study. This chapter evaluates ten single-firm event studies under various Bayesian models. First we provide a brief discussion of the standard event study methodology and the potential invalidity of standard inference that relies on the assumption of normally distributed abnormal returns. Then we discuss the Bayesian perspective and how a Bayesian approach overcomes this potential shortcoming. Our empirical analysis considers ten event studies using data from six firms that have experienced events resulting in a change in their stock prices. We consider Bayesian models that assume stock returns are distributed by a normal distribution and Student t distribution. After considering parametric Bayesian models, we move on to flexible Bayesian models, offering a form of nonparametric Bayesian analysis.

Generally, we offer a valid methodology for conducting single-firm event study inference that does not require a potentially invalid assumption of normality. This allows a court or regulator to argue the effects of an event using valid statistical methods. Our results offer insight into the benefits and implications of using a Bayesian event study, both parametric and flexible, rather than a standard event study.

The posterior distribution determined through the Bayesian approach allows us to

[3]For example, the method for determining damages in securities fraud cases under SEC Rule 10b-5 is not outlined. Therefore, results from inference often are used to determine damages. If the outcome of the event study is based on invalid assumptions, a damages calculation is unreliable. See e.g. Dyl (1999), Cornell and Morgan (1990), and Bhagat and Romano (2002).

conduct Bayesian inference for the event effect in a variety of ways. For each model, we analyze an event effect point estimate and a 95% credible set of the event effect's posterior distribution. The point estimate is determined by the median of the posterior distribution. We find that within the ten events we analyze, this point estimate remains quite similar among the various Bayesian models we consider, including our Bayesian model that replicates the frequentist's ordinary least squares point estimate. This does not imply that the same outcomes of inference result in both the Bayesian and frequentist settings. Because the Bayesian and frequentist perspectives differ, the interpretations from inference are distinct. In particular, these differences may affect how loss causation is established and interpreted. On the other hand, a total damages calculation may not be as affected by these differences in the two approaches.

Damage calculations in securities litigation frequently rely on the ordinary least squares estimate of the event effect. Typically, this point estimate is used to determine a per-share loss. Within each event and for each model we consider, because our resulting event effect medians are quite similar to each another and to the frequentist ordinary least squares point estimate, a calculation of per-share damages would be similar using any of Bayesian point estimates or the frequentist point estimate.[4]

While the event effect's median point estimate might be used to determine total damages, other methods of Bayesian inference likely would be used by the parties to establish loss causation. Establishing loss causation in the frequentist setting relies on hypothesis testing of the event effect. In a Bayesian setting, parties would argue loss causation using the event effect posterior distribution, and determining loss

[4]Within each event, the dissimilarities between the point estimates of each Bayesian model we examine become more economically significant in a damage calculation as the total shares affected increase. For example, the difference in total damages based on a 12 percent point estimate of the event effect loss versus a 12.5 percent estimate of the event effect loss becomes more profound as total shares affected increase.

causation may require more than just using a median point estimate. For example, Bayes factors, due to Jeffreys (1939), is a Bayesian approach to hypothesis testing. Generally, this method of inference involves a ratio of Bayesian models to determine the posterior probability that one of the models is correct (see also Kass and Raftery (1995)). This approach might be used to establish loss causation by comparing a model hypothesizing a nonzero event effect to a model hypothesizing a zero event effect. The Bayes factor is a ratio of the likelihoods of the event effect given the two hypotheses. The posterior odds that the event occurred is the Bayes factor times the prior odds that the event occurred. A posterior odds greater than one indicates that the data supports the hypothesis that the event effect is nonzero over the hypothesis that the event effect is zero. Establishing loss causation would then depend on the interpretation of this posterior odds value. Each party could argue significance of the model supporting the event using a scale developed by Jeffreys (1939), often referred to as "Jeffrey's scale," used for evaluating a posterior odds value.

Another method of Bayesian inference is investigation of the event effect posterior distribution credible set. Within each event, while we find similar event effect posterior medians for each Bayesian model analyzed, the posterior credible sets are more varied. In fact, one Bayesian model for one of our events regarding the case *Apollo Group Inc v. Policemen's Annuity and Benefit Fund of Chicago (2011)* ("Apollo") results in a credible set unlike the other Bayesian models investigated, including the Bayesian model replicating the frequentist results. These differences may affect the determination of loss causation. Additionally, because a credible set might be used to determine a range of per-share damages, varied credible sets could affect the calculation of total damages.[5]

[5]Experts frequently provide a range of damage values that are presented to the court for damage calculations.

A 95% credible set is the interval of event effects for which we have 95% belief that the true event effect lies therein. In the frequentist model, rejection of the null hypothesis that the event effect is zero with a five percent error rate is shown by exclusion of zero from the 95% confidence interval. Under Bayesian analysis, inference of the event effect could oppose inference results determined under frequentist methods. For example, in the aforementioned Apollo case, one Bayesian model we examine results in a credible set that does not contain zero. However, under standard frequentist inference, the researcher fails to reject the null hypothesis that Apollo's event effect is zero. Therefore, additional Bayesian inference could result in the researcher concluding one believes with high probability the event effect is not zero, contrary to the results from standard inference used by the court.

Overall, our results reveal that Bayesian single-firm event studies can overcome the potentially invalid assumption of normally distributed abnormal returns. We also exhibit the possibility that a Bayesian analysis of an event study may result in inference conflicting that of the classical, standard approach. Further, the Bayesian perspective may be attractive in a litigation setting.

2. Overview of the event study methodology

The purpose of an event study is to determine the reaction to a particular event, once public, on a firm's stock returns. Event studies are used in securities litigation to prove materiality and loss causation, as well as determine damages. Fama, Fisher, Jensen, and Roll (1969) and Brown and Warner (1980),(1985) provide the foundations for event study analysis.

The key components of an event study are the event, the event window, and the factors used to identify the event effect. Any combination of market, industry, or firm-specific factors are used to identify the part of a firms's daily stock return affected by an event. An event study hinges on the analysis of "abnormal returns," which is the amount of return for each trading day analyzed that is not explained by the model. Inference is made on the event day abnormal return to determine statistical and economic significance.

The event day is the day or days the event was made public. Announcements of corrective disclosures, stock splits, or mergers are examples of events analyzed in event studies. To complete an event study, the researcher must determine the event window, as well as control period, called the pre-event or post-event period. In the standard event study approach, the pre-event period generally consists of the immediate 100 or 200 days prior to the first event in the event window. Determining the event window involves consideration of the event itself. An event window begins at the first event and continues, at the researcher's discretion, until or even following the last event. For example, an event window might include the day following the announcement, or it might include the thirty days in between a series of events. Multiple events can be jointly or individually analyzed. Securities litigation typically involves few events and often only one.

Upon regressing a firm's daily stock return on a vector of control factors, the estimated residuals determine the abnormal returns for that security. Inference for the event day abnormal return is based on a standard t-test. For valid inference under the standard approach, the data generating process must be normally distributed. However, evidence suggests the distribution of excess returns tends not normally distributed.[6] As previously discussed, even if excess returns are not normally distributed,

[6] See Geweke (2005), Gelbach, Helland, and Klick (2010), and Brown and Warner (1985).

inference based on normality could be valid if the estimated abnormal returns are asymptotically normal. However, in a single-firm and single-event setting like securities fraud litigation, asymptotic normality is beside the point because a central limit theorem cannot apply under such conditions.

Several scholars have developed methods to overcome the potentially invalid inference that arises when assuming abnormal returns are normally distributed. Gelbach, Helland, and Klick (2010) provide evidence that an assumption of normally distributed abnormal returns does not necessarily hold in a single-firm setting. They offer an easily implementable, nonparametric test called the SQ test as an alternative, valid test of classical inference for a single-firm event study. Klick and Sitkoff (2008) provide an alternative method of inference through bootstrapping, as do Ikenberry, Lakonishok, and Vermaelen (1995).

Analyzing abnormal returns through Bayesian statistics, rather than classical methods, offers another alternative. Gelfand and Sfiridis (1996) and Sfiridis and Gelfand (2002) survey the application of Bayesian methods to financial data, and Brav (2000) applies Bayesian methods to long-horizon event studies. None focus on Bayesian analysis of event studies in the single-firm setting.

Zellner (1975) discusses general benefits of the Bayesian approach. At a minimum, analyzing abnormal returns with Bayesian statistics allows the researcher to address the potentially invalid assumption of normally distributed abnormal returns required for standard inference. With the increased ease of Bayesian computation, analyzing abnormal returns from a Bayesian perspective is less costly than it once might have been. Additionally, the Bayesian interpretation is particularly useful in a court setting, where the interested parties care primarily about the effect of a single event.

In a Bayesian setting, since the unknown parameters of the model are treated as random, inference is based on the posterior distribution of the unknown parameters. Generally, for an observed random variable $Z = (z_1, ..., z_e)$ and vector of unknown parameters θ, Bayesian statistics involves determining a posterior distribution of the unknown parameters θ given the vector of observed Z. By Bayes' rule, the conditional posterior density of the unknown parameters is:

$$p(\theta|Z) = \frac{f(Z|\theta)\pi(\theta)}{f(Z)}.$$

The numerator is composed of the density of the data given the unknowns, i.e. the likelihood function, and the marginal density of the unknowns called the prior distribution. The denominator is the marginal density of the data. Since we condition on the observed data, treating the model's parameters as if they are random variables, then the marginal density for the data, $f(Z)$, is a constant. Therefore, the conditional posterior distribution of the unknown parameters, $p(\theta|Z)$, is determined by the prior distribution of the unknown parameters, $\pi(\theta)$, updated by the data and multiplied by some constant. Since the conditional posterior density always will be known up to some constant, we write:

$$p(\theta|Z) \propto f(Z|\theta)\pi(\theta).$$

Depending on the composition of the likelihood and the marginal density, a closed form posterior may not exist. In this case, Markov chain Monte Carlo (MCMC) algorithms may be used to draw from the unknown posterior density. This gives the researcher a sample of draws from the posterior from which to conduct Bayesian inference on the parameters of interest. The posterior distribution's descriptive statistics, a credible set, or a predictive density might be used to conduct Bayesian inference.

2.1. Bayesian versus frequentist interpretation

The difference between the Bayesian and frequentist points of view lies in the interpretation of probability. In the frequentist framework, a probability is interpreted as the frequency of occurrence under repeated trials. However, probability represents a degree of believability in Bayesian statistics.

The frequentist approach assumes a model's parameter has a single, true value, and the data is random. Through multiple random experiments, with some probability the estimated parameter will be the true parameter. To express this probability, one presents a range of values that includes, with some level of "confidence," the true value of the parameter. A 95% confidence interval means one expects the true value of the parameter to fall within that interval 95% of the time a random experiment is conducted to determine the value of the parameter.

The Bayesian perspective, on the other hand, assumes the parameter of interest does not necessarily have one true value. Instead, the parameter's value is determined by a subjective prior probability distribution, based on beliefs (informative or noninformative) about that parameter. Thus, in the Bayesian world, the parameter value is random while the data is fixed, resulting from a single experiment. Then, using the observed data from the experiment, the Bayesian methodology updates this prior belief by the observed data, the combination of which determines a posterior distribution for the parameter(s) of interest. A credible set presents an interval of values from this posterior distribution representing a particular probability of the data, such as 95%. Therefore, under frequentist analysis, one interprets a 95% confidence interval as the interval for which we expect (*ex ante*) 95% of the time to contain the true parameter value. On the contrary, under Bayesian analysis, a 95% credible set is the interval for which we believe (*ex post*) with 95% probability the true parameter lies in that interval.

In a single-firm, single-event study analysis, the Bayesian point of view is not necessarily unreasonable. In court, the frequentist must argue from an underlying belief that repeated random experiments allow one to determine with some probability the true value of the parameter and hence infer particular values for that parameter. However, plaintiffs, defendants, judges, jurors and regulators possibly care less about the probability of that event under realizations over repeated trials and more about the degree of belief regarding the particular event under review. A benefit of the Bayesian approach in this setting is it allows for inference based on data from the experiment of interest. Opponents of the Bayesian methodology, on the other hand, argue that the Bayesian approach requires one to believe not only the prior assumptions about the parameters, but also that one believes with some probability the true parameter value, a random variable, lies in that credible interval.

We proceed with our analysis by first setting up the event study model and describing the data. We consider various Bayesian models to independently analyze ten event studies among six firms. In Section 3.2, we present the classical results using the standard approach. In Sections 4 and 5, our Bayesian analysis considers models based on normal and Student t likelihood functions with varying priors. Finally, in Section 6, we offer a flexible Bayesian model using an empirical distribution function with varying prior distributions. Throughout our analysis, inference of the event effect parameter is based on the posterior median and 95% credible set.

3. Setup of an event study

We employ a market returns model, commonly used in event studies, which assumes a linear relationship between the firm's daily stock return R_s and market factors M_s for each trading day s:

$$R_s = \alpha + \beta M_s + \epsilon_s, \quad s = 1, ..., e-1. \tag{2.1}$$

Multiple covariates could be included, but frequently assumed is a model based on a single explanatory variable, a market returns variable, using a market index such as the NASDAQ or the S&P 500. In our analysis, we assume an event window equal to one day, the event day. This event day occurs at trading day $e = 101$, with a pre-event period of $e - 1 = 100$ trading days.

In the classical event study methodology, ordinary least squares regression over pre-event trading days $s = 1, ..., e-1 = 100$ is used to estimate the parameters of the model, assuming (often implicitly) normally distributed abnormal returns: $\epsilon_s \sim \mathcal{N}(0, \sigma^2)$. The event day abnormal return, ϵ_e, is the event day actual return on the security less the predicted return on the security:

$$\epsilon_e = R_e - (\hat{\alpha} + \hat{\beta} M_e).$$

Standard inference in the classical event study methodology computes a standardized event date abnormal return and uses a t-test to determine statistical significance.

A convenient way to emphasize the effect of the event day is to rewrite the market returns model in equation (2.1) by defining the abnormal return, ϵ_s, as follows:

$$\epsilon_s = \gamma D_s + \mu_s.$$

Variable D_s is defined by an indicator function equal to one on the event day and zero otherwise: $D_s = \mathbf{1}(s = e)$. Thus, on the event date, $\epsilon_e = \gamma + \mu_e$ and all non-event observations yield $\epsilon_s = \mu_s$. Parameter γ is the event effect, or the part of the return from the event day that is not explained by the rest of the model. The market model from equation (2.1) now becomes:

$$R_s = \alpha + \beta M_s + \gamma D_s + \mu_s, \quad s = 1, ..., e. \tag{2.2}$$

In this case, a classical event study estimates the model's parameters using ordinary least squares regression for trading days $s = 1, ..., e$, then calculates the t-statistic for the estimated event effect, $\hat{\gamma}$, to perform standard inference.

As discussed in Gelbach, Helland, and Klick (2010), with exactly one event, the event date fitted abnormal return, $\hat{\mu}_e = R_e - \hat{\alpha} - \hat{\beta} M_e - \hat{\gamma} = 0$, and the ordinary least squares estimates $\hat{\alpha}$ and $\hat{\beta}$ are the same whether the parameters are estimated with or without the event day observation. These two facts imply that the estimated standard error $\hat{\sigma}$ also is the same whether we include the event day observation or not. We use these facts later in our analysis.

3.1. Data and events

We observe daily stock returns and daily market returns originally obtained from the University of Chicago Graduate School of Business Center for Research in Security

Prices (CRSP) via The University of Pennsylvania Wharton Research Data Service (WRDS). For trading days $s = 1, ..., e = 101$, the daily stock return for a particular firm, R_s, is the daily rate of return on a security, calculated as the ratio of the change in the security's closing price to its previous day's closing price.[7] The firm's market return, M_s, is the CRSP value-weighted market return (including dividends) composed of securities from the major indices.

We evaluate ten events occurring among the following six firms: Hershey Corp. (Hershey), Gold Kist, Inc. (Gold Kist), Belo Corp. (Belo), SunPower Corp. (SunPower), Bare Escentuals, Inc. (Bare Escentuals), and Apollo Group, Inc. (Apollo). Five events involving Hershey and Gold Kist are unrelated to securities litigation. Klick and Sitkoff (2008) evaluated the three events in Hershey, which are related to a majority shareholder's sale of stock. For Gold Kist, we evaluate two trading days that experienced extreme changes in the daily return. The remaining five events involving Belo, SunPower, Bare Escentuals and Apollo concern actual events resulting in securities litigation. Of those five events, four were used to prove loss causation and materiality under SEC Rule 10b-5. In the securities litigation case *Apollo Group Inc v. Policemen's Annuity and Benefit Fund of Chicago (2011)*, we analyze an event determined by both parties to be statistically insignificant under the standard approach and therefore not a basis for loss causation. Since this event was dismissed by the court due to statistical insignificance of the event effect, we are interested particularly in whether this event's results from Bayesian inference substantially differ from the results from standard inference used by the court. Section 8: Appendix B provides background information for each event used in our analysis.

[7]Depending on the model, an event study might use log-returns rather than returns.

3.2. Benchmark: classical event study

Using the classical event study methodology described above, Table 2.1 summarizes the results from standard inference. Apollo's event effect is not statistically significant for a one-sided or two-sided test. Thus, the court did not consider this event as part of the plaintiffs' argument for materiality and loss causation. All other event effects are statistically significant at the one percent level of significance for both a one-sided and two-sided test. Accordingly, in the Belo, SunPower, and Bare Escentuals litigation cases, the classical event study provided evidence of loss causation.

Table 2.1. Standard Event Study Estimation Results

Variable	Event Effect	(Std. Error)	(Std. t-statistic)
Hershey event 1	0.2550	0.0125	20.3829***
Hershey event 2	-0.0464	0.0126	-3.6728***
Hershey event 3	-0.1173	0.0125	-9.3803***
Gold Kist event 1	-0.1241	0.0223	-5.5624***
Gold Kist event 2	0.4746	0.0186	25.5615***
Belo event 1	-0.0599	0.0105	-5.7175***
Belo event 2	-0.0330	0.0100	-3.2910***
SunPower	-0.1851	0.0392	-4.7201***
Bare Escentuals	-0.1289	0.0188	-6.8405***
Apollo	-0.0217	0.0197	-1.0988

*** 1% significance level

4. Bayesian analysis under normality

Next, we derive and present analytical results for a Bayesian event study model analogous to the classical event study. This Bayesian model, based on the assumption of normally distributed abnormal returns, generates a mean of the event effect posterior distribution equal to the estimated event effect from ordinary least squares. Then we offer graphical evidence that abnormal returns in our ten event studies in fact are not generated by a normal distribution.

Using market model equation (2.2), for trading day s define data vector $Z_s = (R_s, X_s)$ with $X_s = (1, M_s, D_s)$, and let $\theta = (\alpha, \beta, \gamma)$ and $\sigma^{-2} = \tau$. Assume a flat prior on the location parameter θ and an inverse prior on the scale parameter τ:

$$\pi(\theta|\tau) \propto 1 \quad \text{and} \quad \pi(\tau) \propto \frac{1}{\tau}$$

where by Bayes' rule, $\pi(\theta, \tau) = \pi(\tau)\pi(\theta|\tau)$. Assume

$$\mu_s|X_1 = x_1, ..., X_e = x_e \overset{iid}{\sim} \mathcal{N}(0, \sigma^2), \qquad s = 1, ..., e$$

implying

$$R_s|X_1 = x_1, ..., X_e = x_e, \theta, \sigma \overset{iid}{\sim} \mathcal{N}(x_s'\theta, \sigma^2), \qquad s = 1, ..., e.$$

A well known result is that the conditional distribution for θ is multivariate normal:

$$\theta|\tau, z \sim \mathcal{N}(\hat{\theta}, \sigma^2(X'X)^{-1})$$

where $\hat{\theta}$ is the least squares estimator. The marginal posterior distribution for τ is a scaled chi-square distribution:

$$\tau|z \sim \frac{\chi^2_{(e-k)}}{(e-k)\hat{\sigma}^2}$$

where $k = 3$ total regressors, $e = 101$ trading days, and the estimator for σ^2 is:

$$\hat{\sigma}^2 = \frac{1}{e-k}\sum_{s=1}^{e}(r_s - x_s'\hat{\theta})^2.$$

Because the conditional distribution for θ has a known functional form, we can directly draw from the event effect posterior distribution.[8]

Table 2.2. Normal likelihood: Replication of classical results

	Credible Set			
Event	**2.5%**	**median**	**97.5%**	**(Std Error)**
Hershey event 1	0.2301	0.2549	0.2797	0.0126
Hershey event 2	-0.0716	-0.0465	-0.0216	0.0127
Hershey event 3	-0.1423	-0.1175	-0.0927	0.0126
Gold Kist event 1	-0.1695	-0.1243	-0.0792	0.0230
Gold Kist event 2	0.4363	0.4744	0.5124	0.0194
Belo event 1	-0.0813	-0.0600	-0.0386	0.0109
Belo event 2	-0.0531	-0.0331	-0.0129	0.0103
SunPower	-0.2641	-0.1855	-0.1071	0.0400
Bare Escentuals	-0.1670	-0.1291	-0.0914	0.0193
Apollo	-0.0610	-0.0219	0.0172	0.0200

Comparison of Tables 2.1 and 2.2 reveals that the median of the event effect posterior distribution is nearly equivalent to the ordinary least squares event effect estimate, $\hat{\gamma}$.[9] These results suggest a researcher can perform a single-firm event study using a Bayesian model that replicates the estimated event effect determined under the classical event study methodology. While the interpretation of inference differs between the two frameworks, if abnormal returns are normally distributed, the same valid conclusions can be reached in a Bayesian framework as in the frequentist framework. Under the Bayesian methodology, additional inference based on the event effect posterior distribution other than the median might be used to determine loss

[8] Throughout our analysis of parametric Bayesian models, we use WinBUGS (Bayesian inference Using Gibbs Sampling), a statistical software for Bayesian analysis of statistical models requiring MCMC methods. For consistency, the results in Table 2.2 are based on a WinBUGS model with a normal likelihood, flat priors on the model's location parameters and an inverse prior on the standard deviation.

[9] Computation error accounts for the differences; directly sampling from the known event effect posterior distribution yields a mean of the event effect posterior equivalent to the ordinary least squares estimate of the event effect.

causation, such as Bayes factors described above. The median might be used to determine per-share damages, resulting in a similar damage calculation as under the frequentist approach.

The question remains whether the Bayesian model analyzed above is the correct model. First, we have assumed abnormal returns are normally distributed. Second, we impose no information about the prior distributions of the parameters. Before we investigate models using more informative priors, first we explore the assumption of normally distributed abnormal returns.

Prior research suggests stock returns do not follow a normal distribution. As indicated above, Geweke (2005) discusses the evidence for a Student t distribution of financial data. Gelbach, Helland, and Klick (2010) and Brown and Warner (1985) offer evidence for lack of normality.

Figures 2.1, 2.2, and 2.3 for firms SunPower, Gold Kist event 1, and Hershey event 3, respectively, provide plots of the estimated kernel densities of fitted standardized abnormal returns with a superimposed standard normal density. These three density graphs, as well as the density graphs for the other seven events, reveal that for each event we consider, the estimated fitted abnormal returns generally are not generated exactly by a normal distribution. For most of the events studied, the estimated standardized abnormal returns densities have smaller spread than the standard normal density.

In each of the three graphs, the area to the left of the fifth quantile of the standard normal density ends left of the area representing the fifth quantile of the estimated standardized abnormal returns density. This implies that by using the standard approach, the researcher will reject the null hypothesis less than five percent of the time, therefore concluding abnormal returns are statistically significant less often than one should under valid standard inference. Since our graphical analysis suggests that stock returns are not generated from a normal distribution, we proceed with Bayesian models relevant to our beliefs about abnormal returns that do not rely on an assumption of normality.

5. Bayesian analysis under Student t models

Based on prior research that the distribution of stock returns has fatter tails than a normal distribution, we consider a Student t likelihood function under varying priors. An additional parameter of the Student t distribution is the degrees of freedom. We can fix the degrees of freedom, place a prior distribution on the degrees of freedom, or use an estimate of the degrees of freedom.

We implement three Bayesian models with fixed degrees of freedom 3, 10, and 40. Recall, a Student t distribution with infinite degrees of freedom is equivalent to a standard normal distribution. Results, provided in Tables 2.4 – 2.13, reveal credible sets similar to that of the normal model with the same prior distributions on the remaining parameters.[10] With a Student t likelihood, the event effect posterior cannot be derived analytically. Therefore, we use Gibbs sampling to draw from the posterior distribution. As expected, as the degrees of freedom increase, a Student t model yields results closely resembling that of a model based on normality.

Within each event, each Student t model's event effect posterior median is nearly identical to the normal model's event effect posterior median. The larger difference in magnitudes are among the bounds of the credible sets. For every event except

[10] We also consider a prior distribution on the degrees of freedom based on estimating the degrees of freedom via maximum likelihood for a model with assumed Student t errors. However, since varying the degrees of freedom for the events we analyze does not yield substantial changes in inference, a model based on a maximum likelihood estimate of the degrees of freedom will add nothing.

SunPower and Apollo, the bounds of each model's credible set become tighter as the degrees of freedom increase. That is, the upper and lower bounds of the 95% credible sets decrease and increase, respectively, as we move closer to a Student t model that approximates a normal model. Given the redistribution of mass from the center to the tails of the Student t distribution as the degrees of freedom decrease, these results are reasonable.

To apply the results of Bayesian inference to a calculation of damages, a court could use the event effect posterior median as a measure of each share's loss due to the fraudulent behavior of the firm. In the case of Student t models, the value would differ only slightly from the event effect estimate determined under classical inference. However, total damages also depend on the total number of shares affected by the violation. Therefore, the more shares affected, the more sensitive the total damages calculation is to differences among the event effect posterior medians for various Bayesian models.

Thus far, our analysis reveals two main results. First, we show a Bayesian model assuming normally distributed abnormal returns yields a mean of the event effect posterior distribution identical to (accounting for computation error) the event effect's ordinary least squares estimate from a classical event study. However, both of these classical and Bayesian single-firm event study models assume normally distributed abnormal returns, which may fail to hold.

If the abnormal returns in fact come from a normal distribution, one advantage of the Bayesian approach over the classical approach is the determination of a posterior distribution of the event effect, rather than just a point estimate or confidence interval. One may use the event effect posterior distribution to conduct various types of inference. Additionally, the Bayesian interpretation of probability may be attractive in a litigation environment.

If the abnormal returns fail to come from a normal distribution, as often believed in the finance literature and as indicated by our evaluation of ten events, then the Bayesian methodology allows for an event study that does not rely on normally distributed abnormal returns. A belief that stock returns often follow at Student t distribution provides the basis for the second Bayesian model we consider. Our second main result reveals that for the ten events we analyze, the posterior median and 95% credible set for a Student t model with known degrees of freedom are nearly identical across degrees of freedom and nearly equivalent to the normal model. Accordingly, for the events we consider, if one believes abnormal returns follow a Student t distribution, a Bayesian analysis results in the median of the event effect posterior distribution nearly the same as the ordinary least squares estimate under the classical setting. However, in the Bayesian Student t model, inference does not rely on an assumption of normality, as does classical inference.

In practice, parametric Bayesian event study analysis of a single firm offers a method of statistical inference that does not require a potentially invalid assumption of normally distributed excess returns. Even if a Bayesian event study model assuming Student t abnormal returns results in valid inference, in a litigation setting parties may disagree on the distribution of abnormal returns. If abnormal returns do not exactly follow a Student t model, then even under Bayesian analysis, inference that plays a large role in determining the outcome and damages of a case will be inaccurate. Rather than assume a possibly incorrect distribution of abnormal returns, one could apply a nonparametric Bayesian analysis using the exact distribution of the pre-event abnormal returns data. In the next section, we investigate the benefits of a flexible Bayesian model.

6. Bayesian analysis under flexible models

Thus far we have assumed a parametric form on the distribution of the data. For the ten events we examine, since we graphically illustrate that abnormal returns do not exactly follow a normal distribution, our analysis indicates the results from standard inference are invalid. The distribution of abnormal returns might more closely resemble some other known distribution besides the normal and Student t. However, rather than investigate other potential parametric Bayesian models, we do away with assumptions on the form for the abnormal returns distribution and explore a flexible Bayesian event study model.

Gelbach, Helland, and Klick (2010) develop a nonparametric test for determining statistical significance of a single-firm event study's estimated event effect. Rather than assume a distributional form for the abnormal returns, they use the empirical distribution function of the pre-event estimated residuals as an estimate of the unknown distribution of abnormal returns. Their "SQ test" for a five percent level of statistical significance compares the fifth most-negative pre-event estimated abnormal return to the ordinary least squares event effect estimate. If the estimated event effect is less than or equal to this fifth percentile of the pre-event sample, then the researcher rejects the null hypothesis that the event had no effect on that firm's daily stock return.

Like Gelbach, Helland, and Klick (2010), our flexible Bayesian analysis also estimates the cumulative distribution function of the abnormal returns by using the empirical distribution function for the pre-event abnormal returns. However, we proceed in a Bayesian rather than frequentist setting by using the empirical distribution function for the abnormal returns to form a conditional posterior distribution of the data.

A flexible model overcomes invalid inference that can result from an event study model, classical or Bayesian, based on an assumption of normally distributed abnormal returns. However, a flexible model also has practical benefits. Agreement among expert witnesses regarding the correct model for a given situation is unlikely and therefore administratively costly for courts. Inference based on a flexible model has the added benefit of universal use by the courts. Assumptions for the most appropriate prior distribution of the event effect would be argued by each party. In our flexible model analysis, we explore informative and noninformative prior distributions and examine the implications of the results, as they relate to securities litigation.

First, we develop our flexible Bayesian model. Begin with the market model equation (2.2):

$$R_s = \alpha + \beta M_s + \gamma D_s + \mu_s, \quad s = 1, ..., e = 101.$$

Since the pre-event period is relatively large, we ignore estimation error in the ordinary least squares estimates $\hat{\alpha}$ and $\hat{\beta}$ and treat them as if they are the true parameters. We let α and β be the value of $\hat{\alpha}$ and $\hat{\beta}$. Therefore, uncertainty lies only in the event effect parameter, γ. On event date e, the market model equation reduces to:

$$\begin{aligned} Z_e &= R_e - \alpha - \beta M_e \\ &= \gamma + \mu_e. \end{aligned} \tag{2.3}$$

For prior distribution $\pi(\gamma)$ and likelihood function $f_Z(z_e|\gamma)$, we desire to draw from the posterior distribution for event effect γ given the data:

$$\begin{aligned} p_\Gamma(\gamma|Z = z_e) &= \frac{f_Z(z_e|\gamma)\pi_\Gamma(\gamma)}{f_Z(z_e)} \\ &\propto f_Z(z_e|\gamma)\pi_\Gamma(\gamma). \end{aligned} \tag{2.4}$$

We ignore estimation error in $\hat{\mu}_s$, the fitted residual on day s from regressing pre-event daily returns on market returns, thereby assuming μ_s is known, equivalent to $\hat{\mu}_s$. Let $\hat{F}_\mu(t)$ be the empirical distribution function of μ_s:

$$\hat{F}_\mu(t) = \frac{1}{e-1}\sum_{i=1}^{e-1} \mathbf{1}(\mu_i \leq t).$$

We take $\hat{F}_\mu(t)$ to be the distribution of μ_s in order to form the likelihood function. The conditional distribution of the data, $f_Z(z_e|\gamma)$, is transformed into a likelihood function in terms of abnormal returns μ in the following way:

$$\begin{aligned} f_Z(z_e|\gamma) &= Pr(Z = z_e|\gamma) \\ &= Pr(\gamma + \mu = z_e|\gamma) \\ &= Pr(\mu = (z_e - \gamma)|\gamma) \\ &= \hat{f}_\mu(z_e - \gamma). \end{aligned} \tag{2.5}$$

Typically, the likelihood function in a Bayesian model includes parameters for which the researcher places prior distributions. The empirical distribution function used as an estimate for the likelihood function acts as if the parameters of the likelihood have been determined outside of the Bayesian model. Our estimated, discrete likelihood function, $f_Z(z_e|\gamma) = \hat{f}_\mu(z_e - \gamma)$, yields the following posterior distribution for the event effect given the assumed-known parameter values α and β:

$$\begin{aligned} p_\Gamma(\gamma|Z=z_e) &\propto f_Z(z_e|\gamma)\pi(\gamma) \\ &= \hat{f}_\mu(z_e-\gamma)\pi(\gamma). \end{aligned} \tag{2.6}$$

For various event effect prior distributions, $\pi(\gamma)$, we use Monte Carlo simulation to draw from the posterior distribution. We begin with a flat prior, then consider more informative priors. Finally, we compare the results from five flexible models under various prior distributions on the event effect to the normal model analyzed in Section 4. For each model, the event effect posterior median and 95% credible set are provided in Tables 2.3, 2.14 – 2.22 and discussed in Section 6.3.

6.1. Flexible model with a flat prior

With a flat prior distribution on the event effect, $\pi(\gamma) \propto 1$, the event effect posterior distribution is determined by direct draws from $\hat{f}_\mu(z_e-\gamma)$. For $j = 1, ..., J$, we draw $\mu^{(j)}$ from $\hat{F}_\mu(t)$, then form $\gamma^{(j)} = z_e - \mu^{(j)}$. This yields J draws from the event effect posterior distribution $p_\Gamma(\gamma|Z=z_e)$ provided in equation (2.6). From these draws, we then determine the 95% credible set and median.

6.2. Flexible models with informative priors

In a litigation setting, parties could use *ex ante* beliefs about the event effect abnormal return to assume a flexible model with informative priors. With informative priors, since we cannot determine a functional form of the posterior distribution, we cannot directly draw from the posterior distribution. We employ importance sampling to draw from the event effect posterior distribution.

6.2.1. Importance sampling

Importance sampling is a method of estimating a desired statistic from a posterior distribution when no closed-form posterior distribution exists. Weighted draws from a "source density," also known as an importance function, are used in place of the target posterior density. The source density, which should include the support of the posterior density, is a convenient density from which we can directly draw. We weight these draws to correct for any bias. The better the approximation of the source density to the target posterior, the more reliable the importance sampling estimates.

Suppressing the conditioning on the data, define $p_\Gamma(\gamma)$ as the target posterior density and $p_I(\gamma)$ as the source density from which we can directly draw. Define the weighting function as the ratio of the target posterior to the source density:

$$w(\gamma) = \frac{p_\Gamma(\gamma)}{p_I(\gamma)}.$$

Generally, for some function of the event effect, $h(\gamma)$, rewrite the following expectation so it is taken over source density $p_I(\gamma)$, from which we can directly draw, rather than target density $p_\Gamma(\gamma)$, from which we cannot directly draw:

$$\begin{aligned} E(h(\gamma)) &= \frac{\int_\Gamma h(\gamma) p_\Gamma(\gamma) d\gamma}{\int_\Gamma p_\Gamma(\gamma) d\gamma} \\ &= \frac{\int_\Gamma \frac{h(\gamma) p_\Gamma(\gamma)}{p_I(\gamma)} p_I(\gamma) d\gamma}{\int_\Gamma \frac{p_\Gamma(\gamma)}{p_I(\gamma)} p_I(\gamma) d\gamma} \\ &= \frac{\int_\Gamma [h(\gamma) w(\gamma)] p_I(\gamma) d\gamma}{\int_\Gamma w(\gamma) p_I(\gamma) d\gamma}. \end{aligned} \tag{2.7}$$

For draws $j = 1, ..., J$ from the source density, use equation (2.7) to approximate $E(h(\gamma))$ by

$$\frac{\frac{1}{J}\sum_j w(\gamma^{(j)})h(\gamma^{(j)})}{\frac{1}{J}\sum_j w(\gamma^{(j)})}.$$

In our flexible Bayesian model, we use the empirical distribution function as our importance function. Substituting the defined target posterior $p_\Gamma(\gamma)$ from above with the desired conditional posterior $p_\Gamma(\gamma|Z = z_e)$ from equation (2.6), the weighting function simplifies to the prior distribution of event effect γ:

$$\begin{aligned} w(\gamma) &= \frac{p_\Gamma(\gamma)}{p_I(\gamma)} \\ &= \frac{\hat{f}_\mu(z_e - \gamma)\pi(\gamma)}{\hat{f}_\mu(z_e - \gamma)} \\ &= \pi(\gamma). \end{aligned}$$

We use importance sampling to generate draws from the target posterior density, then determine the median and a 95% credible set from those draws. For $j = 1, ..., J$, first we draw $\gamma^{(j)}$ from the source density: $\gamma^{(j)} = z_e - \mu^{(j)}$, where $\mu^{(j)}$ is drawn from $\hat{F}_\mu(t)$. Next we determine a weighted percentile for the median and bounds of the 95% credible set. To do this, we order the draws, calculate the weighting vector (based on the assumed prior distribution), then normalize the weighting vector. From the normalized weighting vector, we calculate the cumulative sum of the normalized weights. Using the cumulative sum vector, for the values nearest and not less than 0.025, 0.50, and 0.975, we determine each of the corresponding three ranks. For the draws from the source density, the corresponding order statistic for each of these three ranks determines each desired percentile, yielding a 95% credible set and median from draws from the target event effect posterior distribution.

6.2.2. Informative priors selection

In a model assuming an informative prior distribution for the event effect, an appropriate prior would be argued by each party. For example, the most elementary prior information known about the event effect posterior distribution is the limitation on its support. Using information known about the event, we can place a known bounded interval $[a, b]$ on the flat prior distribution's support to impose an informative prior distribution in our model, $\gamma \sim \mathcal{U}(a, b)$. For $z_e = \gamma + \mu$, this is equivalent to bounding draws $\mu^{(j)}$ from the empirical distribution function in the following way:

$$a \leq \gamma \leq b \quad \iff \quad (z_e - a) \geq (\mu = z_e - \gamma) \geq (z_e - b).$$

In cases where one believes the event had a negative impact on the firm's stock prices (abnormal returns are nonpositive), the most extreme prior belief is that the event effect is bounded by $[-1, 0]$ since the daily return can decrease by no more than 100 percent and no less than zero percent.

For a more lenient bound on an event effect that could arguably reduce or increase the firm's value, we consider a prior belief that an event decreases returns by no more than 100 percent, and increases returns by no more than ten fold, a generous upper bound. This results in a uniform event effect prior distribution with support [-1,10]. These two supports on the uniform distribution may be so generous that these two informative priors contribute to the updating of our posterior no more than the flat prior, resulting in the same posterior distribution as that determined under the flexible, flat prior model. Nonetheless, both parties likely would agree to these two prior distributions:

$$\gamma \sim \mathcal{U}(-1, 0) \quad \text{and} \quad \gamma \sim \mathcal{U}(-1, 10).$$

For our third flexible model with an informative prior, since many argue stock returns are generated from a Student t distribution rather than a normal distribution, we consider the prior belief that the event effect is generated by a Student t distribution with four degrees of freedom:

$$\gamma \sim t(4).$$

Finally, we also consider a prior belief that the event effect is generated by Laplace distribution, also known as a double exponential distribution, which is more peaked than a Student t distribution:

$$p(\gamma) = \frac{1}{2} \exp(|\gamma - \eta|)$$

where η is the location parameter, which we assume known at 0.25. While one could assume prior distributions on the degrees of freedom for the Student t prior and the location parameter for the Laplace prior, we are interested in an investigation of the general methodology. Thus, we keep the Bayesian models simple by assuming known values for these parameters.

6.3. Discussion of results

We compare the results from the five flexible Bayesian models to the normal model from Section 4. For convenience, credible sets and medians for the Apollo event are provided below in Table 2.3, while results for the remaining nine events are provided in Tables 2.14 – 2.22.

Table 2.3. Apollo: Flexible models

Likelihood	**Prior on γ**	**2.5%**	**median**	**97.5%**
Normal	Flat	-0.0610	-0.0219	0.0172
EDF	Flat	-0.0513	-0.0222	0.0247
EDF	$\mathcal{U}(-1,0)$	-0.0513	-0.0240	-0.0023
EDF	$\mathcal{U}(-1,10)$	-0.0513	-0.0222	0.0247
EDF	t(4)	-0.0513	-0.0222	0.0247
EDF	Laplace$(0,0.25)$	-0.0492	-0.0218	0.0247

First, notice that the medians for each posterior distribution for all flexible models within an event are nearly the same as each other and nearly the same as the median under the normal model. Recall that for all Bayesian models considered in our analysis, two events, Hershey event 1 (Tables 2.4 and 2.14) and Gold Kist event 2 (Tables 2.8 and 2.18), result in a positive event effect posterior median, while the other eight events result in a negative event effect posterior median. For Hershey event 1 and Gold Kist event 2, restricting the flat prior distribution of the event effect to a support on $[-1,0]$ results in no importance sampling estimate. Under importance sampling, zero weight is given to each draw from their respective empirical distribution functions because no draws from their empirical distribution functions are nonpositive.

Within all events, the $\mathcal{U}(-1,10)$ and $t(4)$ flexible models result in exactly the same credible sets and medians as under the flexible model with a flat event effect prior distribution. Furthermore, for all events except Hershey 1 and Gold Kist 2 (the two events yielding no posterior distribution under this flexible model), as well as Apollo, the $\mathcal{U}(-1,0)$ flexible model yields exactly the same credible sets and medians within each event as three other flexible models: the flat prior model, the $\mathcal{U}(-1,10)$ model, and the $t(4)$ model. Therefore, for all events except Apollo, only the double exponential flexible model within each event yields credible sets not identical to the remaining flexible models' credible sets.

A closer comparison of the identical credible sets to those under the normal model reveals subtle differences in the credible sets *among* events. For example, for SunPower (Table 2.21), the four identical flexible models result in larger event effect medians (corresponding to smaller negative event effects) than the normal model median. SunPower's credible sets' bounds are tighter (larger lower bound and smaller upper bound) than the bounds of the normal model.

However, for Bare Escentuals (Table 2.22) and Apollo (Table 2.3), the respective identical flexible models' posterior medians are smaller (corresponding to a larger negative event effect) than the normal model's posterior median. Also, the bounds of the identical flexible models' credible sets are shifted right when compared to the bounds of the normal model, towards larger event effect bounds (corresponding to smaller negative event effect bounds). Therefore, in the case of Bare Escentuals, for example, the defendant could offer the identical 95% credible sets from the flexible models as evidence of a range of values for which one believes with 95% probability the true event effect lies in that interval. Compared to the credible set from the normal model, these ranges of event effects are less negative than those of the normal model, resulting in an interval of lower per-share damages.

With respect to the double exponential model, for all three of Hershey's events (Tables 2.14 - 2.16), both of Belo's events (Tables 2.19 and 2.20), and Gold Kist event 2 (Table 2.18), the upper and lower bounds of the credible sets are identical to the identical models' credible sets' bounds. However, for all of the events, the medians of the double exponential are not identical to the normal or other flexible models' medians. For all events, the posterior medians under a double exponential prior are larger than the identical flexible models' median, resulting in either a smaller negative event effect or a larger positive event effect. For all of the securities litigation cases we analyze, if damage calculations were based on the posterior median or the credible set intervals, then the double exponential model is attractive to the plaintiff for some events and attractive to the defendant for other events. The bounds of the double exponential credible sets compared to the bounds of the normal model's credible sets vary among events: for Belo event 1 and Bare Escentuals the bounds are smaller in absolute value (hence a range of smaller event effects) than the normal model; for Belo event 2 the bounds are larger in absolute value (hence a range of larger event effects) than the normal model; and for SunPower the bounds are tighter. Regardless, the attractiveness of the model also depends on the magnitude of the differences between the credible sets' values. As explained above, the greater the total shares affected by the event, the potentially larger effect any differences among results have on a total damages calculation.

For the Apollo event, a flexible model with a uniform event effect prior distribution on [-1,0] results in a tighter credible set than that of the normal model. The upper bound of this model's credible set is a negative event effect, rather than a positive event effect like the normal model, the Student t models (Table 2.13), and the remaining flexible models.

Since this result is unlike any of our other results, we examine the quantile-quantile (q-q) plots for Apollo and Belo 2 (Figure 2.4), for comparison. Both q-q plots graph the quantiles of the distribution of abnormal returns against the quantiles for the standard normal distribution to determine if the data is generated by a standard normal. Overall, the Belo event 2 abnormal returns fit the standard normal much closer than the Apollo abnormal returns. For Apollo, the points distant from the standard normal quantile are much further away than those on the Belo event 2 q-q plot. These distinct points on the Apollo q-q plot indicate the distribution of abnormal returns is possibly skewed more to the right, but at least has longer tails than under a standard normal distribution. If we interpret the Apollo q-q plot as an indication of right skewness, we can justify our importance sampling results for Apollo's $\mathcal{U}(-1, 0)$ flexible model. If the right tail of the distribution of abnormal returns is longer than the left tail, we tend to draw from a more negative pool of abnormal returns. Combining this with the fact that we have restricted our event effect prior distribution to a nonpositive support, the resulting nonpositive credible set is not unreasonable.

This distinctive result from the Apollo flexible model offers an example of how a Bayesian analysis of a single-firm event study can yield results that may conflict with the results under standard, classical inference. Coincidentally, this is the one event in our study dismissed by the court due to the results from a classical event study's standard inference.

Under the classical methodology, the 95% confidence interval contains zero, resulting in a failure to reject the null hypothesis that the estimated event effect is zero. Under a flexible Bayesian model, our 95% credible set excludes zero. Therefore, one might conclude from a Bayesian point of view that we believe with 95% probability the true event effect lies in that interval, which does not contain zero. However, one cannot necessarily conclude the event effect is nonzero. Fortunately, because the Bayesian methodology results in an entire posterior distribution for the event effect, other methods of Bayesian inference might yield further evidence of a zero or nonzero event effect for the Apollo case, which may or may not be in line with the court's decision not to consider the event.

7. Conclusion

We provide insight into the use of Bayesian analysis of single-firm, single-event studies as a valid methodology for determining inference on the effects of an event on a firm's share or bond price. We show how the classical methodology is replicated in a Bayesian setting. Then we compare posterior means of the event effect from various Bayesian models to the estimated event effect under the frequentist approach. Finally, we analyze a nonparametric Bayesian model using the empirical distribution function and various prior distributions.

Within each event, for all Bayesian models considered, we find the medians of the event effect posterior distributions are similar. Differences among the credible sets and medians might be substantial, depending on the court's use of the inference results. If inference is used to determine a per-share damage calculation, then the varied results between the models could be nontrivial as the total number shares affected increase.

While our inference from a Bayesian single-firm event study might not yield vastly different results from that under classical inference, results from Bayesian inference do not require an assumption of normality. The Bayesian approach offers valid inference not guaranteed by the classical, standard approach. Because the courts rely on event studies to determine materiality, loss causation, and damages, inference must be based on valid theory.

Our results under flexible Bayesian models show the bounds of credible sets within an event vary more than the posterior median. Furthermore, analysis of various flexible models reveals additional insight into the possibility of concluding inference on the event effect posterior distribution that conflicts with the results of standard inference under a classical single-firm event study. As Apollo's flexible model under the $\mathcal{U}(-1, 0)$ event effect prior reveals, zero is not in the 95% credible set. This suggests the possibility that with further analysis of the event effect posterior, such as through Bayes factors, one might conclude a result for loss causation opposing that of the court's to disregard this event due statistical insignificance under the classical standard (and potentially invalid) approach.

We apply Bayesian statistics to event studies of single firms and single events to evaluate the benefits and compare results from inference under the standard, classical methodology. With particular focus on determining loss causation and damages in securities litigation, we highlight differences and potential discrepancies in inference under the standard methodology and the Bayesian methodology. Future work includes investigating other methods of Bayesian inference for determining loss causation and damages, as well as examining other settings that apply single-firm event studies to determine if a Bayesian analysis offers further insight.

8. Appendix B: Background on the events

We obtained information regarding the four securities cases from the Stanford Law School Securities Class Action Clearinghouse (2011) website. Information about the Hershey Corporation is based on Klick and Sitkoff (2008).

First, in *Apollo Group Inc v. Policemen's Annuity and Benefit Fund of Chicago (2011)*, a corrective disclosure announcement on September 9, 2004, resulted in an event effect that was not statistically significant at the court's standard five percent level of significance. Therefore, this event was not considered by the court. On September 8, 2004, Apollo Group's stock ended the day at $82.07. On the day of the announcement, the stock ended the day at $80.50. Over the next three trading days, the stock remained near this level, at $80.43 on September 10, $80.63 on September 13, and $80.09 on September 14. Therefore, the event study presented by the parties assumed a four-day event window. We assuming a one-day event window on the day of the announcement, September 9, 2004.

Most event studies not resulting in statistical significance of the event effect never reach the court and therefore are not public record. As indicated above, this event provides the rare opportunity to investigate, with respect to an actual legal dispute, whether the court might not have dismissed the effects of this event if Bayesian inference, rather than potentially invalid standard inference, had been applied. Therefore, valid Bayesian statistical methods might show proof of loss causation, which the court might otherwise dismiss under potentially invalid statistical inference. This event also allows us to analyze the extent to which the conclusions from standard inference, resulting in the court's disregard of this event, resemble the conclusions from valid, Bayesian inference. [11]

[11] The Apollo case eventually reached the U.S. Supreme Court and was decided in favor of the

The other three cases we examine are standard securities litigation cases. In *Fener v. Belo Corp. (2008)*, the plaintiffs alleged that the Belo Corp. intentionally overstated the circulation of its principal newspaper, the *Dallas Morning News*, to increase advertising sales. The overstated circulation estimates also were reported to investors, inflating the firm's financials. We separately consider two events evaluated during litigation. On March 9, 2004, Belo announced a 2.5 percent and 3.5 percent decrease in the future circulation of its daily and Sunday newspapers, respectively (Belo event 2). By the end of the trading day, the Belo stock price decreased from \$28.57 to \$27.58. In a press release Belo issued the evening of August 5, 2004, they admitted an internal investigation found suspect behavior regarding their circulation practices (Belo event 1). Other announcements included the resignation of the head of circulation and intent to refund advertisers the amount overcharged. Around this same time, the plaintiffs alleged Belo admitted it overstated its circulation forecast by 1.5 percent and 5 percent for its daily and Sunday newspapers, respectively. By the end of the next trading day, August 6, 2004, Belo's stock price decreased from \$23.21 to \$21.55.

In *In re Bare Escentuals, Inc. Securities Litigation (2010)*, we analyze an earnings announcement from August 1, 2007, in which the stock price decreased approximately 13 percent, from \$28.30 on August 1, 2007, to \$24.75 on August 2, 2007. Bare Escentuals, Inc. participates in developing, advertising and selling cosmetics under brand names such as bareMinerals. The firm allegedly failed to disclose unfavorable information about the firm's financial condition, such as the decline in sales after a new infomercial campaign.

defendants in March 2011. A second event, publication of a third-party report, which resulted in a statistically signifiant event effect, became the focus of the case. Because the report was not a disclosure by the firm, the case turned on what constitutes a disclosure, per SEC rules.

Finally, in *Sunpower Corp. Systems v. Sunlink Corp. (2009)*, we evaluate an event occurring on the evening of November 16, 2009, when SunPower filed a report with the SEC stating that financial statements (annual reports, and previously reported financial results) are not reliable. By the end of the next trading day, SunPower's stock price decreased by 19 percent, from $27.23 to $22.19. The plaintiffs alleged that SunPower's behavior resulted in investors trading at artificially inflated stock prices.

The three Hershey Corporation events, while not evaluated in securities litigation, are the basis of Klick and Sitkoff (2008). They analyze the effects of an announcement event and two resulting court actions involving the selling of a majority of stock in the Hershey Corporation. On July 25, 2002, *The Wall Street Journal* published a story announcing the intent of the Milton Hershey School Trust (Trust), a super shareholder of Hershey, to sell its Hershey shares (Hershey event 1). On September 4, 2002, the Pennsylvania Attorney General was granted an injunction by a trial judge, preventing the Trust from selling its 30 percent share (75 percent voting right) in Hershey. Finally, on September 18, 2002, the Trust publicly announced its abandonment of the sale (Hershey event 3). The event window lasted from the day of the sale announcement, July 25, 2002, to the day of the sale abandonment, September 18, 2002. However, we analyze returns for each of the three major daily events.

To offer additional variation in the events and daily return changes examined in our analysis, we chose two trading days for which Gold Kist, Corp, experienced extreme changes in daily returns. The first occurred on June 23, 2005, (Gold Kist event 1) when the price decreased nearly 13 percent, from $23.21 to $20.36. The second occurred on August 21, 2006, (Gold Kist event 2) when the price increased 47.5 percent, from $12.93 to $19.02.

TABLE 2.4. Hershey event 1: Normal and Student t likelihoods, flat priors

Likelihood	**2.5%**	**median**	**97.5%**
Normal = Student t(∞)	0.2301	0.2549	0.2797
Student t(40)	0.2288	0.2550	0.2814
Student t(10)	0.2281	0.2552	0.2816
Student t(3)	0.2266	0.2553	0.2843

TABLE 2.5. Hershey event 2: Normal and Student t likelihoods, flat priors

Likelihood	**2.5%**	**median**	**97.5%**
Normal = Student t(∞)	-0.0716	-0.0465	-0.0216
Student t(40)	-0.0729	-0.0466	-0.0200
Student t(10)	-0.0737	-0.0463	-0.0196
Student t(3)	-0.0746	-0.0458	-0.0176

TABLE 2.6. Hershey event 3: Normal and Student t likelihoods, flat priors

Likelihood	**2.5%**	**median**	**97.5%**
Normal = Student t(∞)	-0.1423	-0.1175	-0.0927
Student t(40)	-0.1436	-0.1173	-0.0910
Student t(10)	-0.1442	-0.1172	-0.0908
Student t(3)	-0.1454	-0.1171	-0.0889

Table 2.7. Gold Kist event 1: Normal and Student t likelihoods, flat priors

Likelihood	**2.5%**	**median**	**97.5%**
Normal = Student t(∞)	-0.1695	-0.1243	-0.0792
Student t(40)	-0.1698	-0.1240	-0.0781
Student t(10)	-0.1698	-0.1237	-0.0787
Student t(3)	-0.1724	-0.1230	-0.0742

Table 2.8. Gold Kist event 2: Normal and Student t likelihoods, flat priors

Likelihood	**2.5%**	**median**	**97.5%**
Normal = Student t(∞)	0.4363	0.4744	0.5124
Student t(40)	0.4356	0.4746	0.5138
Student t(10)	0.4352	0.4750	0.5139
Student t(3)	0.4312	0.4759	0.5199

Table 2.9. Belo event 1: Normal and Student t likelihoods, flat priors

Likelihood	**2.5%**	**median**	**97.5%**
Normal = Student t(∞)	-0.0813	-0.0600	-0.0386
Student t(40)	-0.0831	-0.0599	-0.0365
Student t(10)	-0.0835	-0.05965	-0.0362
Student t(3)	-0.0844	-0.05969	-0.0351

Table 2.10. Belo event 2: Normal and Student t likelihoods, flat priors

Likelihood	**2.5%**	**median**	**97.5%**
Normal = Student t(∞)	-0.0531	-0.0331	-0.0130
Student t(40)	-0.0549	-0.0328	-0.0105
Student t(10)	-0.0556	-0.0326	-0.0100
Student t(3)	-0.0561	-0.0324	-0.0083

TABLE 2.11. SunPower: Normal and Student t likelihoods, flat priors

Likelihood	**2.5%**	**median**	**97.5%**
Normal = Student t(∞)	-0.2641	-0.1855	-0.1071
Student t(40)	-0.2476	-0.1837	-0.1195
Student t(10)	-0.2403	-0.1833	-0.1272
Student t(3)	-0.2467	-0.1830	-0.1207

TABLE 2.12. Bare Escentuals: Normal and Student t likelihoods, flat priors

Likelihood	**2.5%**	**median**	**97.5%**
Normal = Student t(∞)	-0.1670	-0.1291	-0.0914
Student t(40)	-0.1683	-0.1294	-0.0904
Student t(10)	-0.1702	-0.1297	-0.0897
Student t(3)	-0.1791	-0.1308	-0.0832

TABLE 2.13. Apollo: Normal and Student t likelihoods, flat priors

Likelihood	**2.5%**	**median**	**97.5%**
Normal = Student t(∞)	-0.0610	-0.0219	0.0172
Student t(40)	-0.0609	-0.0218	0.0176
Student t(10)	-0.0603	-0.0223	0.0152
Student t(3)	-0.0648	-0.0229	0.0194

Table 2.14. Hershey event 1: Flexible models

Likelihood	**Prior on γ**	**2.5%**	**median**	**97.5%**
Normal	Flat	0.2301	0.2549	0.2797
EDF	Flat	0.229552	0.256030	0.278408
EDF	$\mathcal{U}(-1, 10)$	0.229552	0.256030	0.278408
EDF	$t(4)$	0.229552	0.256030	0.278408
EDF	$Laplace(0, 0.25)$	0.229552	0.255610	0.278408

Table 2.15. Hershey event 2: Flexible models

Likelihood	**Prior on γ**	**2.5%**	**median**	**97.5%**
Normal	Flat	-0.0716	-0.0465	-0.0216
EDF	Flat	-0.071838	-0.045361	-0.022983
EDF	$\mathcal{U}(-1, 0)$	-0.071838	-0.045361	-0.022983
EDF	$\mathcal{U}(-1, 10)$	-0.071838	-0.045361	-0.022983
EDF	$t(4)$	-0.071838	-0.045361	-0.022983
EDF	$Laplace(0, 0.25)$	-0.071838	-0.045153	-0.022983

Table 2.16. Hershey event 3: Flexible models

Likelihood	**Prior on γ**	**2.5%**	**median**	**97.5%**
Normal	Flat	-0.1423	-0.1175	-0.0927
EDF	Flat	-0.142802	-0.116325	-0.093947
EDF	$\mathcal{U}(-1, 0)$	-0.142802	-0.116325	-0.093947
EDF	$\mathcal{U}(-1, 10)$	-0.142802	-0.116325	-0.093947
EDF	$t(4)$	-0.142802	-0.116325	-0.093947
EDF	$Laplace(0, 0.25)$	-0.142802	-0.116284	-0.093947

TABLE 2.17. Gold Kist event 1: Flexible models

Likelihood	**Prior on γ**	**2.5%**	**median**	**97.5%**
Normal	Flat	-0.1695	-0.1243	-0.0792
EDF	Flat	-0.174999	-0.121107	-0.083043
EDF	$\mathcal{U}(-1, 0)$	-0.174999	-0.121107	-0.083043
EDF	$\mathcal{U}(-1, 10)$	-0.174999	-0.121107	-0.083043
EDF	$t(4)$	-0.174999	-0.121107	-0.083043
EDF	$Laplace(0, 0.25)$	-0.169713	-0.120388	-0.069328

TABLE 2.18. Gold Kist event 2: Flexible models

Likelihood	**Prior on γ**	**2.5%**	**median**	**97.5%**
Normal	Flat	0.4363	0.4744	0.5124
EDF	Flat	0.435334	0.476281	0.507891
EDF	$\mathcal{U}(-1, 10)$	0.435334	0.476281	0.507891
EDF	$t(4)$	0.435334	0.476281	0.507891
EDF	$Laplace(0, 0.25)$	0.435334	0.475221	0.507891

TABLE 2.19. Belo event 1: Flexible models

Likelihood	**Prior on γ**	**2.5%**	**median**	**97.5%**
Normal	Flat	-0.0813	-0.0600	-0.0386
EDF	Flat	-0.080328	-0.060796	-0.036159
EDF	$\mathcal{U}(-1, 0)$	-0.080328	-0.060796	-0.036159
EDF	$\mathcal{U}(-1, 10)$	-0.080328	-0.060796	-0.036159
EDF	$t(4)$	-0.080328	-0.060796	-0.036159
EDF	$Laplace(0, 0.25)$	-0.080328	-0.060206	-0.036159

TABLE 2.20. Belo event 2: Flexible models

Likelihood	**Prior on γ**	**2.5%**	**median**	**97.5%**
Normal	Flat	-0.0531	-0.0331	-0.0129
EDF	Flat	-0.054243	-0.031717	-0.018666
EDF	$\mathcal{U}(-1, 0)$	-0.054243	-0.031717	-0.018666
EDF	$\mathcal{U}(-1, 10)$	-0.054243	-0.031717	-0.018666
EDF	$t(4)$	-0.054243	-0.031717	-0.018666
EDF	$Laplace(0, 0.25)$	-0.054243	-0.031547	-0.018666

TABLE 2.21. SunPower: Flexible models

Likelihood	**Prior on γ**	**2.5%**	**median**	**97.5%**
Normal	Flat	-0.2641	-0.1855	-0.1071
EDF	Flat	-0.234337	-0.182270	-0.124443
EDF	$\mathcal{U}(-1, 0)$	-0.234337	-0.182270	-0.124443
EDF	$\mathcal{U}(-1, 10)$	-0.234337	-0.182270	-0.124443
EDF	$t(4)$	-0.234337	-0.182270	-0.124443
EDF	$Laplace(0, 0.25)$	-0.233610	-0.180345	-0.122045

TABLE 2.22. Bare Escentuals: Flexible models

Likelihood	**Prior on γ**	**2.5%**	**median**	**97.5%**
Normal	Flat	-0.1670	-0.1291	-0.0914
EDF	Flat	-0.157901	-0.130905	-0.083865
EDF	$\mathcal{U}(-1, 0)$	-0.157901	-0.130905	-0.083865
EDF	$\mathcal{U}(-1, 10)$	-0.157901	-0.130905	-0.083865
EDF	$t(4)$	-0.157901	-0.130905	-0.083865
EDF	$Laplace(0, 0.25)$	-0.156183	-0.129333	-0.080807

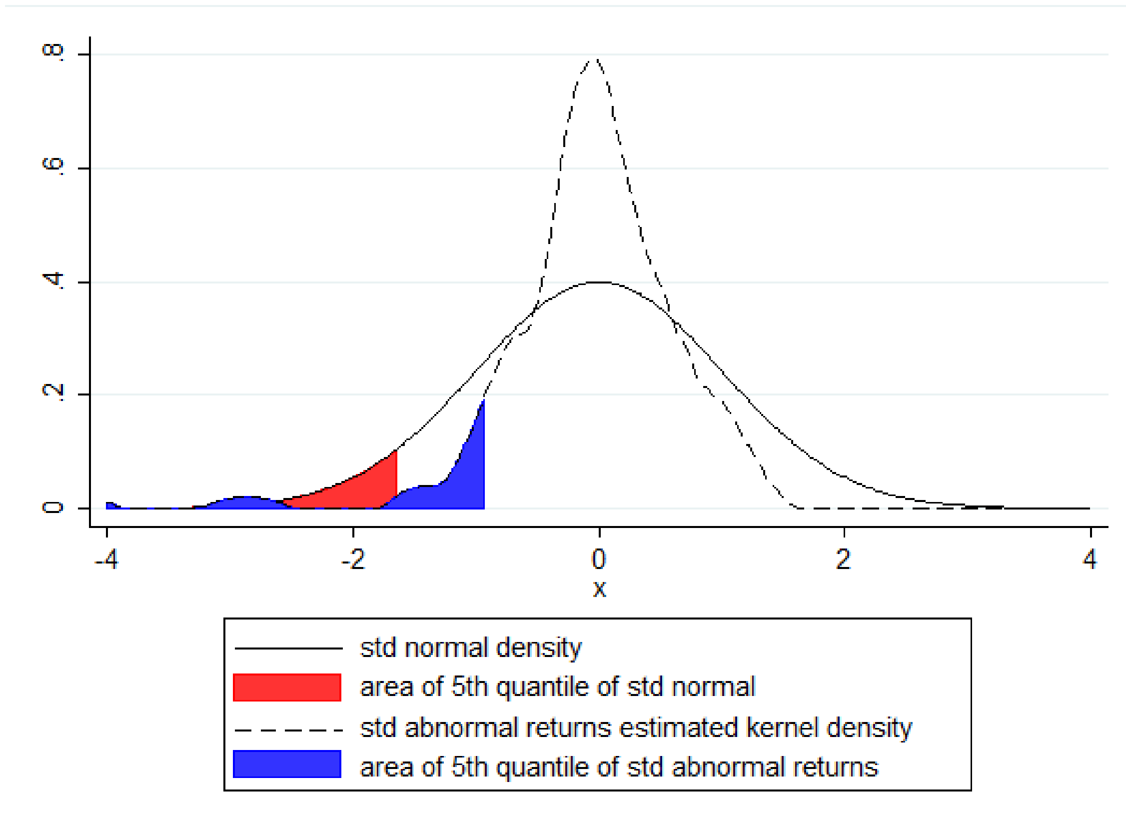

FIGURE 2.1. SunPower: Std Normal Density vs Est Std Abnormal Returns Density

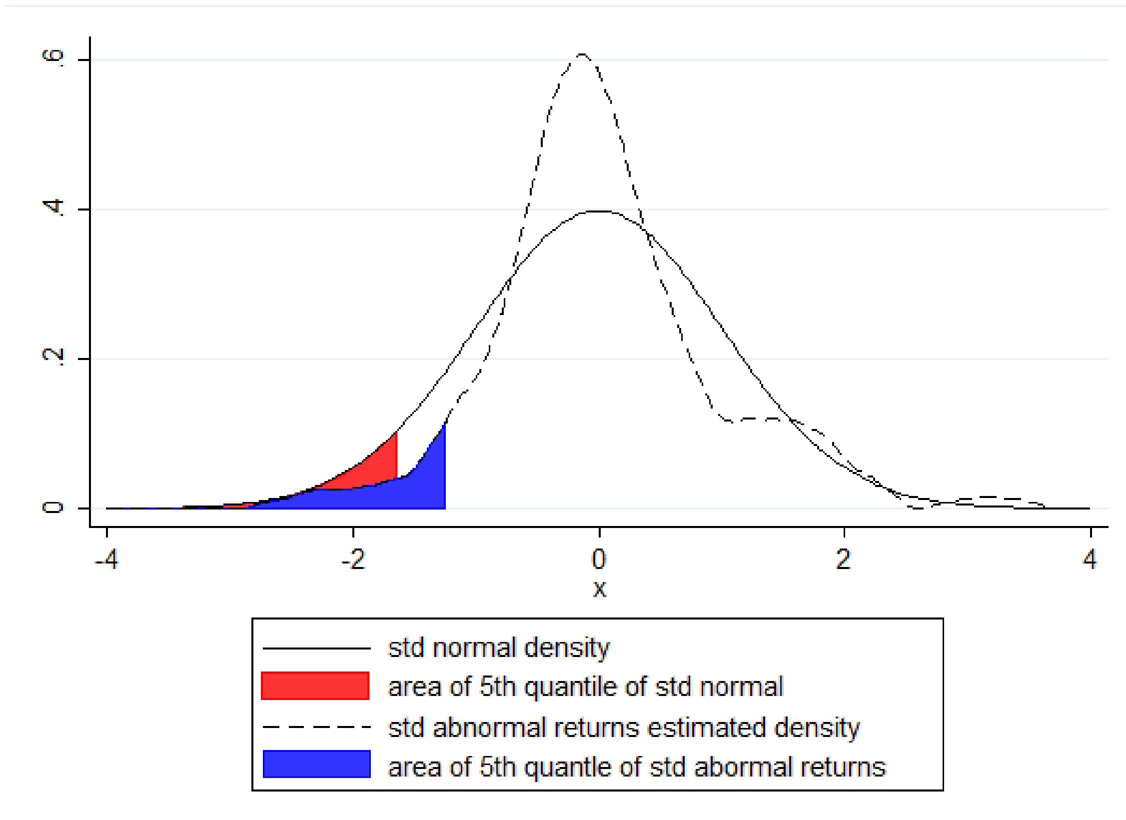

FIGURE 2.2. Gold Kist event 1: Std Normal Density vs Est Std Abnormal Returns Density

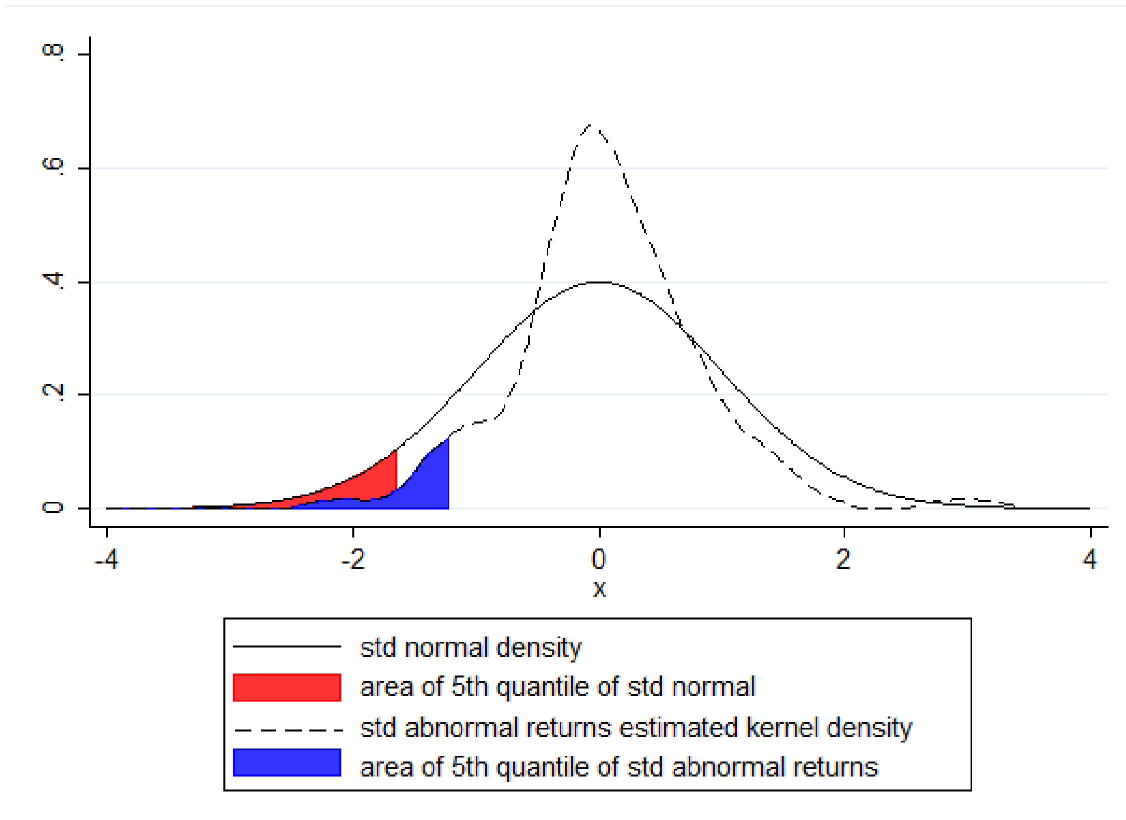

Figure 2.3. Hershey event 3: Std Normal Density vs Est Std Abnormal Returns Density

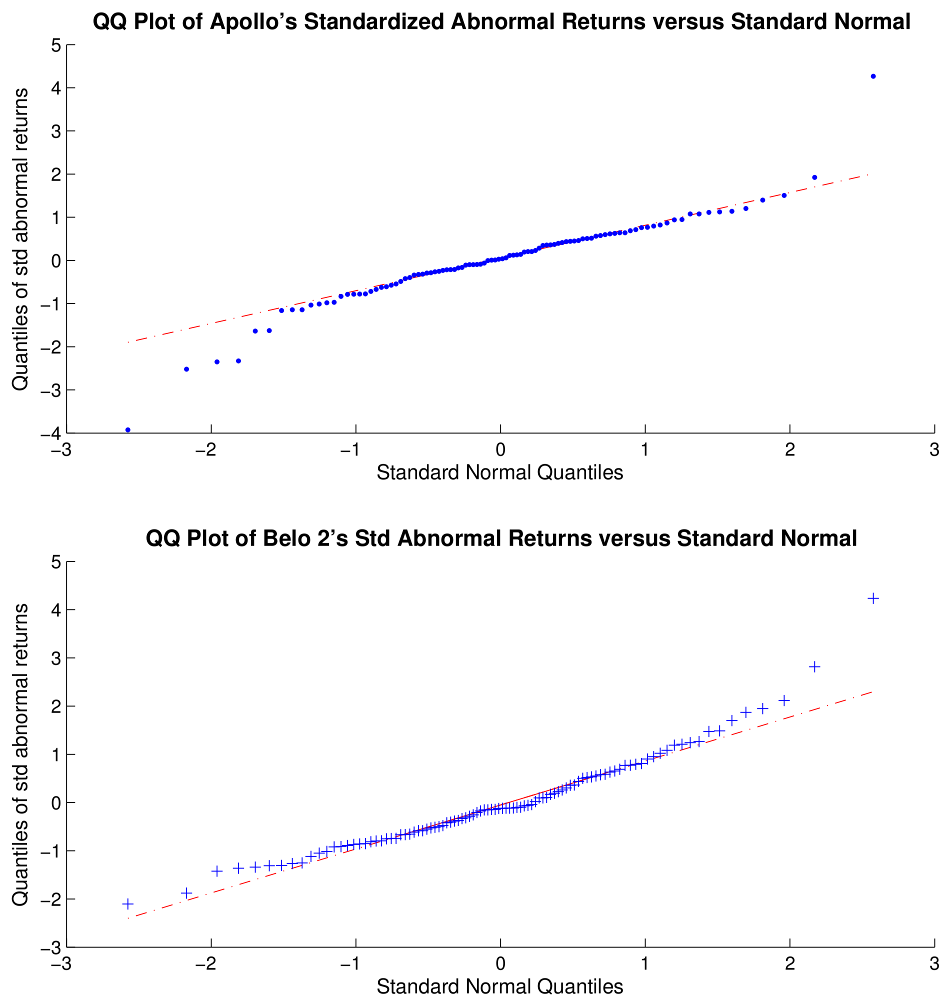

FIGURE 2.4. Apollo and Belo 2: Q-Q Plots of Standardized Abnormal Returns vs Standard Normal

Chapter 3

Firm and Consumer Preferences for Information: An Application to Pharmaceutical Advertising

1. Introduction

In 2009, Bayer Healthcare Pharmaceuticals, producer of oral contraceptive, Yaz®, launched a $20 million advertising campaign, peaking television viewers' attention with the following:

> *"You may have seen some Yaz commercials recently that were not clear. The FDA wants us to correct a few points in those ads..."*[1]

This unusual campaign was the result of a settlement agreement between Bayer and the Food and Drug Administration (FDA), along with 27 states. The FDA believed two previous Yaz® commercials downplayed health risks associated with the use of Yaz® while overstating possible side benefits. Bayer avoided potentially harsher sanctions by agreeing to clarify the side effects and benefits in a corrective advertising campaign. They also agreed to FDA approval of all Yaz® commercials for the next six years.[2]

Bayer's advertising of Yaz® is a form of direct-to-consumer advertising (DTCA), or mass media pharmaceutical marketing disseminated to the general public, as opposed to only physicians. *Ex ante*, the end-user may have little or no information

[1] See, for example, *The New York Times* article by Singer (2009). At the date of this manuscript, at least one of the advertisements was viewable via *YouTube.com* (see Bayer Pharmaceuticals (Creator), posted February 26, 2009).

[2] See the Warning Letter to Bayer, U.S. Food and Drug Administration (2008a) and the Illinois Attorney General (2009) announcement.

about a particular medical condition and treatment. DTCA's potential to *increase* consumer valuations for a product typically motivates a firm's use of DTCA. However, DTCA also can *disperse* consumer valuations, depending on the type of information provided. When DTCA presents a drug's risks and benefits, such as those for Yaz®, some consumers increase, while others might decrease their initial valuations. In the case of the Yaz® broadcasts, the FDA grew concerned about consumers' ability to determine the true quality of the drug, as the risks and benefits were misrepresented. If DTCA obscures the true quality of a product, consumers are less likely to correctly ascertain their true valuations for the good.

The Yaz® advertisement dispute between Bayer and the FDA highlights a firm's strategic provision of profit-maximizing levels of information, constrained by a regulator's goal to maximize consumer welfare. This chapter theoretically examines the implications for dispersing consumer valuations through this strategic behavior. Motivating our results through pharmaceutical DTCA provides justification for and against regulating DTCA.

Our analysis employs a framework developed by Johnson and Myatt (2006) for analyzing rotations in demand. New or updated product information typically motivates a shift in a demand curve. A firm's marketing strategy may affect consumers' willingness-to-pay, but the resulting effect on the shape of demand is not necessarily a shift. If advertising unanimously persuades consumers, resulting in a homogeneous change in their valuations, then demand shifts. However, new information, obtained through marketing or advertising, can affect a consumer's sensitivity to price changes (e.g, Halmenschlager, Mantovani, and Troege (2011)). If new information persuades some while dissuading others, then consumers' adjustments to the information is not homogeneous; the distribution of consumer valuations changes shape, thereby resulting in *at least* a rotation in demand. Johnson and Myatt's framework is summarized in Sections 2 and 3.2.

The theoretical model, analysis, and application presented herein is an extension of Hawkins and Lazzati (2011). We build on Johnson and Myatt (2006) by further analyzing the effects of an exogenously determined information dispersion parameter, such as advertising, in an oligopoly market for a homogeneous product. We examine the effects of this parameter on firms' profits and consumer welfare, as well as determine how production costs and entry into the market affect optimal levels of advertising.

Johnson and Myatt (2006) show that under the assumptions of a constant elasticity of demand, firms' profits are U-shaped in the dispersion parameter. We show this result holds similarly for consumers, but also that the marginal benefits from increasing information are higher for consumers than for firms. Therefore, firms and consumers prefer either the maximal or minimal level of information provision, with consumers always preferring more.

We can reinterpret the scale of maximal to minimal information provision in terms of *real* versus *hype* advertising. A rotation in the demand curve increases with real information and is minimized with hype information. Firms offering more information pursue a *niche* market strategy, while minimal information corresponds to a *mass* market strategy. Section 2.1 characterizes DTCA in terms of real and hype information under both niche market and mass market advertising strategies.

Since consumers always prefer more information than firms, preferences for advertising do not necessarily align. For example, suppose firms prefer to provide hype DTCA and production costs are low. This generally is the case once generic firms enter the pharmaceutical market. Then consumers might prefer more real product information, such as additional information about risks and negative side effects. Our results offer justification for regulation of DTCA. However, we show as firms enter the market or as the cost of production decreases, preferences for advertising align to the minimum levels. In fact, there exists some level of entry or some level of production cost at which firms and consumers prefer the same extremal (maximal or minimal) level of advertising. Therefore, under certain market conditions, preferences may align.

The remainder of this chapter explains the demand rotation framework of Johnson and Myatt (2006), then presents and analyzes a model of Cournot competition under constant elasticity of demand. We then motivate our model of firm and consumer preferences for advertising in terms of pharmaceutical DTCA. We conclude with a discussion of how our results may apply to DTCA.

2. Background on demand rotation

Determinants of demand accompany any standard explanation of the theory of supply and demand. However, the theory behind demand rotation has only recently gained ground and still is seldom addressed (Graves and Sexton (2006)). Empirical research has shed light on the causes and implications of demand rotations. For example, Ackerberg (2001) measures informative and prestige advertising effects, Meyerhoefer and Zuvekas (2008) estimate the shape of demand, and Zheng, Kinnucan, and Kaiser (2010) measure and test rotations of demand that stem from pharmaceutical advertising.

Theoretical analyses of demand rotations lacks a common framework. Kinnucan and Zheng (2004) note the theoretical result that a rotation in demand does not necessarily imply demand elasticity shifts. Aislabie and Tisdell (1988) address demand rotations in monopolies and consider the effects of ‘bandwagon’ or ‘snob’ advertising strategies. Graves and Sexton (2006) briefly comment on income changes as an underlying cause of a rotation and highlight the importance of distinguishing shifts from rotations. Johnson and Myatt (2006) offer a general framework for demand rotations, motivated by the provision of information. Their framework analyzes the effects of changes in an exogenous information parameter on monopoly profits. As they indicate, and as addressed below, their framework applied to a Cournot oligopoly setting is complicated by strategic effects arising from competition. Hence, our analysis focuses on extending their framework to the Cournot oligopoly setting by analyzing a constant elasticity demand specification.

The level of information chosen by firms ranges from low (hype) to high (real). This level determines the degree of information dispersion that results in a possible rotation of demand. Firms’ preferences for dispersion depend on the relationship between the equilibrium quantity and the quantity at which dispersion causes a rotation in demand (the rotation point). In our model, following the provision of information that results in some level of rotation, firms compete in Cournot quantities.

As shown in Figure 3.1, an increase in some advertising level a from a' to a'' causes a clockwise rotation in inverse demand. To the left of the quantity $Z^{\dagger}$, at which the clockwise rotation occurs, consumer valuations increase, while to the right of this rotation point, consumer valuations decrease. If the marginal consumer’s willingness-to-pay *decreases* upon receiving the product information, then the firm’s profits will be lower with the increase in dispersion than their profits would be otherwise. In this case, firms prefer as little dispersion as possible. When minimal dispersion of consumer valuations is optimal for firms, the advertising strategy consists of provid-

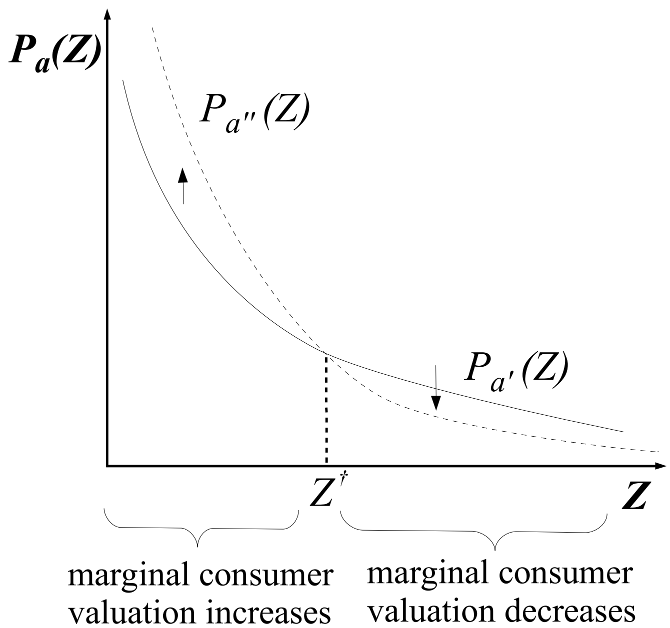

FIGURE 3.1. Clockwise demand rotation

ing less real product information, i.e supplying the mass market. In other words, firms prefer that consumers remain more alike in their updating of valuations for the product. In the extreme, hype advertising does little to disperse valuations, providing the most basic product information.

However, firms welcome an increase in dispersion if the marginal consumer's willingness-to-pay *increases* due to dispersion. In Figure 3.1, this occurs for quantities to the left of the rotation point $Z^{\dagger}$. In this case, dispersion of consumers' valuations yields higher profit as valuations have increased for the consumers willing to purchase the good. Even though the increase in information might result in a loss of some consumers willing to buy the good, those remaining in the market have increased their valuations. In such situations, we say firms pursue a niche market by providing a high level of real product information.

Regardless of firms' preferences for advertising levels, regulation constrains firms to particular forms of information provision. DTCA is one form of information regulated in the U.S. by both the FDA and the Federal Trade Commission (FTC). Next, we motivate demand rotation through types of DTCA and the regulation thereof.

2.1. Demand rotation through DTCA

Since the 1980s, DTCA has been a growing form of pharmaceutical advertising. Shortly after its introduction, and until 1997, the FDA required that broadcast advertisements include the same information provided in writing to physicians (similar to that provided on a one-page DTCA print advertisement, such as in a magazine), making broadcast advertising prohibitively expensive (Ling, Berndt, and Kyle (2002)). An FDA study in 1985 revealed that consumers desired more information about pharmaceuticals than had been offered by the industry. In 1997, the FDA issued updated regulatory guidelines for broadcasted pharmaceutical DTCA, weakening prior guidelines. The new guidelines, for example, now allow pharmaceutical advertisers to state the name of the drug *without* fully disclosing all effects of the drug. Bayer's Yaz® broadcasts are a form of this advertising.

Since 1997, pharmaceutical DTCA has greatly increased in the U.S. The pharmaceutical industry's expenditures on DTCA increased from $791 million in 1996 to $2.5 billion in 2000. In 2000, three years after the new guidelines took effect, television advertising accounted for 57% of total mass media prescription drug advertising costs, increasing 27.3% between 1999 and 2000 (NIHCM Foundation (2001)). However, nearly all other countries continue to prohibit DTCA (Pines (1999), Palumbo (2002)).

While the FDA has addressed consumers' desires for more pharmaceutical drug information, its new guidelines obscure their requirement of "a fair balance between benefit and risk information" in pharmaceutical broadcast advertising.[3] Prior to 1997, advertising required a maximal level of *real* information, ensuring consumers received very detailed information on risks and negative effects of a drug. However, now the FDA's regulation of DTCA consists of three types of advertising, described below. Each type regulates a firm's advertising behavior when firms choose to offer hype or real information. Furthermore, each type of DTCA advertising corresponds to firms serving a mass or niche market in terms of how the particular DTCA changes consumer valuations for the product.

Under one type of DTCA, the FDA allows firms to provide consumers the product name and its intended use, but it also requires a balanced presentation of the drug's benefits and risks in any given advertisement. This FDA mandate results in information provided to consumers about product risks that otherwise might not be provided. Information about the medical condition and treatment might increase consumer valuations, causing a shift in demand. However, consumers could respond to detailed information about risks and negative effects in various ways, thereby dispersing consumer valuations and causing a rotation in demand for the drug. An advertising campaign in which consumers' valuations disperse upon receipt of more detailed product information can be considered a niche market strategy. Dispersing consumer valuations through detailed information might decrease the mass of consumers who are willing and able to purchase the good, but for the remaining niche, valuations for the good have increased.

A second form of DTCA, "reminder advertisements," is exempt from the FDA requirement on disclosure of risks and negative effects as long as the intended use is withheld from the advertisement. A third form of DTCA, "help-seeking adver-

[3] See U.S. Code 21 C.F.R. §202.

tisements," is not subject to FDA regulation when a particular drug name is not mentioned in the advertisement. These two forms of DTCA could be considered mass market strategies, focusing on as little dispersion of consumer valuations as possible by providing basic product information. In such a case, demand rotations would be minimized as compared to a niche market strategy.

Opponents of DTCA often claim that regulation, or lack thereof, does not require enough information in advertisements about the negative effects and possible risks of the drug (see, e.g, Iizuka (2004)). Surveys by the FDA in 1999 and 2002 showed that DTCA does not equally convey information about risks and benefits. As with the sanction on Bayer for its Yaz® advertising campaign, the FDA seeks to determine if firms supply information according to their requirements for each type of DTCA permitted. Our results show that firms might not choose to provide as much information as preferred by consumers, in which case, allowing hype information might decrease welfare, and regulation may be justified. We proceed with our analysis and results.

3. Framework and analysis

To begin, we set up our model of Cournot oligopoly competition with an exogenous advertising parameter. Before analyzing our model, we briefly describe the theoretical framework of demand rotation developed in Johnson and Myatt (2006).

3.1. Model setup

Assume the market is populated by a unit mass of consumers who decide whether to acquire a good. A consumer is willing to pay up to θ for the product, where θ is drawn from distribution $F_a = 1 - \exp(\mu(a)/a)\theta^{-1/a}$ with support on $\Theta_a = (\exp(\mu(a)), \infty)$.

Here, $a \in A = [\underline{a}, \overline{a}] \subset [0,1]$, captures firms' advertising efforts and indexes a family of distribution functions.

Given a price p, a fraction $Z = 1 - F_a(p) = \exp(\mu(a)/a)p^{-1/a}$ of consumers purchase the good. If Z units are to be sold, then the market clearing price satisfies $P_a(Z) = F_a^{-1}(1-Z)$ or $P_a(Z) = \exp(\mu(a))Z^{-a}$ for all $Z \in [0,1]$. This yields an inverse elasticity of demand $\eta_a(Z) = a$ for all Z. Differentiating the latter with respect to a yields:

$$\frac{\partial \eta_a(Z)}{\partial a} = 1 > 0.$$

Therefore, a also indexes a family of inverse demand curves in terms of inverse elasticity of demand.[4]

In the supply side of the market, $n > 1$ identical firms compete in quantities, *à la* Cournot, taking a as given. We assume each firm faces constant marginal cost $c > 0$. We can think of this model as a case of co-opetition where firms are rivalrous in terms of production but coordinate with respect to product information provided to consumers.

3.2. Demand dispersion and rotation ordering

Johnson and Myatt (2006) develop a framework to evaluate the role of information (advertising parameter a in our model) in shifting and/or rotating demand. We describe their *rotation ordering* and how it applies to our framework.

Advertising might change consumer valuations in a number of ways. For example, an increase in a might yield a decrease in $F_a(\theta)$ for all θ. In this case, the family of valuation distributions, $\{F_a(\theta)\}$, is ordered by first-order stochastic dominance. This

[4] Amir and Lazzati (2010) use the dual of this condition (see assumption $A5$) to study the effects of market structure on industry performance in a Cournot game with network effects.

translates into a homogeneous change in consumer valuations that causes the inverse demand curve to shift. Some types of advertising, such as hype advertising where little real information is provided to consumers, might have this effect on demand. However, as described above, consumers might react dissimilarly to advertising, causing some valuations to increase, while others decrease. Under such a heterogeneous change in consumer valuations, a first-order stochastic ordering of valuation distributions no longer applies.

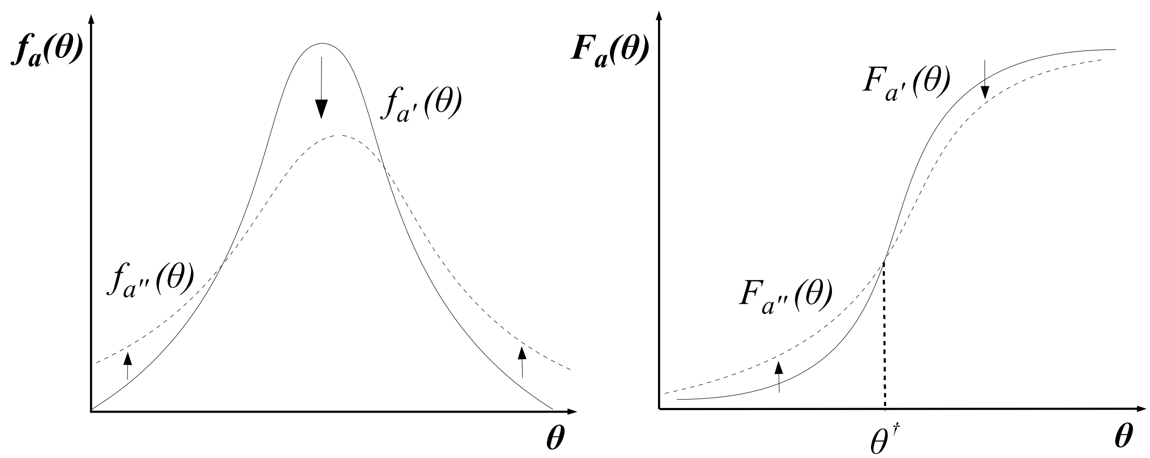

FIGURE 3.2. Spread (and possible shift) in $f_a(\theta)$; Clockwise rotation in $F_a(\theta)$

Johnson and Myatt present a stochastic ordering similar to second-order stochastic dominance. With their *rotation ordering*, provided in Definition 1, an increase in advertising redistributes the mass of consumer valuations from the center to the tails of the density function, $f_a(\theta)$. This spread (and possible shift) in $f_a(\theta)$ corresponds to a clockwise rotation of the distribution of valuations $F_a(\theta)$ around some point of rotation, $\theta^{\dagger}$. This corresponds to the inverse demand function $P_a(z)$ rotating around quantity $Z^{\dagger}$ as shown in Figure 3.1 for $a' < a''$. As shown in Figure 3.2, for $a' < a''$ density $f_a(\theta)$ spreads and distribution $F_a(\theta)$ rotates.[5] Therefore, a rotation ordering

[5]Johnson and Myatt (2006) include a graph similar to that of Figure 3.2.

satisfies the single-crossing property for the cumulative distribution function: with an increase in advertising, all valuations to the left of the rotation point increase and all valuations to the right of the rotation point decrease.

Definition 1. ***Rotation Ordered*** (Johnson and Myatt (2006))*: A local change in a leads to a rotation of* $F_a(\theta)$ *if for some* $\theta_a^\dagger$ *and each* $\theta \in (\underline{\theta}_a, \bar{\theta}_a)$,

$$\theta \gtrless \theta_a^\dagger \Leftrightarrow \frac{\partial F_a(\theta)}{\partial a} \lesseqgtr 0.$$

If this holds for all a, then $\{F_a(\theta)\}$ *is ordered by a sequence of rotations. Equivalently,*

$$Z \gtrless Z_a^\dagger \Leftrightarrow \frac{\partial P_a(Z)}{\partial a} \lesseqgtr 0$$

where $Z_a^\dagger = 1 - F_a(\theta_a^\dagger)$*, so the family of inverse demand curves* $\{P_a(Z)\}$ *is ordered by a sequence of rotations.*

Johnson and Myatt largely focus on a monopoly analysis. Their rotation ordering places no restriction on how the slope of demand changes in response to advertising. However, extending their analysis to a competitive model requires that we know how equilibrium quantity responds to advertising. Due to strategic effects, the equilibrium quantity under Cournot competition requires one to sign the change in the slope of inverse demand to an increase in advertising. Since this is indeterminate, generally, we cannot analyze a general competitive model without additional restrictions to their framework. Section 6: Appendix C sets up the general framework under Cournot oligopoly competition where strategic interaction among firms results in an ambiguous analysis of the effects of advertising. Therefore, we extend Johnson and Myatt's analysis of Cournot oligopoly competition under constant elasticity of demand to examine the alignment of firm and consumer preferences for advertising.

Under a constant price elasticity of demand assumption, it is straightforward to apply the definition of rotation ordering to show how advertising leads to a clockwise rotation around the following quantity: $Z^{\dagger}(a) = \exp(\mu'(a))$ if $\mu'(a) \leq 0$. That is, the market price increases when $Z < Z^{\dagger}(a)$ and decreases when $Z > Z^{\dagger}(a)$. This description therefore captures the informational role of advertising: after receiving the new piece of information, some consumers are willing to pay more, while others decrease their valuations.[6]

3.3. Effects of advertising on firms and consumers

From a methodological perspective, this analysis relies on lattice programming, or supermodular optimization. Amir (2005) provides an overview of the theory of supermodular games with marked emphasis on accessibility. His survey covers all theorems invoked here.

Under the specification explained in Section 3.1, the oligopolistic competition has a unique equilibrium, and all relevant functions can be easily calculated. Our first result provides an expression for the equilibrium aggregate output and market price.

Proposition 2. *At the unique equilibrium, the aggregate level of production and the market price satisfy:*

$$Z(a,n,c) = \left[\frac{(n-a)\exp(\mu(a))}{nc}\right]^{1/a} \quad \textit{and} \quad P(a,n,c) = \frac{nc}{(n-a)}.$$

It follows from Proposition 2 that the market price increases with advertising.

[6]When $\mu'(a) \geq 0$ advertising shifts the inverse demand function right, so there is no rotation point.

However, the effect of a on total equilibrium quantity $Z(a,n,c)$ is ambiguous. Thus, profits need not necessarily increase with more information, as discussed below.

Proposition 3 reflects equilibrium firm profits and consumer surplus.

Proposition 3. *At the unique equilibrium, Cournot per-firm profits and consumer surplus satisfy:*

$$\pi(a,n,c) = \frac{ac}{n(n-a)} Z(a,n,c) \quad \textit{and} \quad S(a,n,c) = \frac{nac}{(n-a)(1-a)} Z(a,n,c).$$

To evaluate the preferences of firms and consumers regarding advertising, define

$$a_\pi(n,c) = \arg\max\{\pi(a,n,c) : a \in A\} \quad \textit{and} \quad a_S(n,c) = \arg\max\{S(a,n,c) : a \in A\}$$

so that $a_\pi(n,c)$ and $a_S(n,c)$ are the (sets of) most preferred levels of advertising for firms and consumers, respectively. Proposition 4 shows if firms determine that a higher level of a is profitable, then consumers similarly find that a higher level of a increases consumer surplus. Therefore, consumers' most preferred level of advertising is always higher than the level most preferred by firms.

Proposition 4. *The marginal benefits of increasing a satisfy, for all $a, a' \in A$ with $a > a'$,*

$$\pi(a,n,c) - \pi(a',n,c) \geq 0 \implies S(a,n,c) - S(a',n,c) > 0.$$

Thus, every element in $a_S(n,c)$ is higher than all the elements in $a_\pi(n,c)$.

The next proposition states that when $\mu''(a) \geq 0$, both firms and consumers prefer either $\overline{a}$ or $\underline{a}$, but not necessarily the same extreme of A.

Proposition 5. *If $\mu''(a) \geq 0$, profits are strictly convex in a and consumer surplus is quasiconvex in a. Thus, $a_\pi(n,c) \in \{\underline{a}, \overline{a}\}$ and (generically) $a_S(n,c) \in \{\underline{a}, \overline{a}\}$.*[7]

Thus, when $\mu''(a) \geq 0$, the conflict of interest between firms and consumers regarding advertising can be quite severe. In particular, if we combine the last two results, it is possible that while firms would like to coordinate in the lowest level of advertising, $\underline{a}$, consumers would rather the maximum level of information, $\overline{a}$.

While advertising preferences do not necessarily align, Proposition 6 shows that exogenous entry by firms in the market pushes the optimal level of advertising down for *both* firms and consumers.

Proposition 6. *Firms' profits and consumer surplus are strictly log-submodular in a and n:*

$$\frac{\partial^2 \ln \pi(a,n,c)}{\partial a \partial n} < 0 \quad \textit{and} \quad \frac{\partial^2 \ln S(a,n,c)}{\partial a \partial n} < 0.$$

Thus, every selection from $a_\pi(n,c)$ or $a_S(n,c)$ is decreasing in the number of firms in the market, n.

In a similar way, the next result shows that firm and consumer preferences for advertising increase with technological improvements, i.e. a decrease in c.

Proposition 7. *Firms' profits and consumer surplus are strictly log-supermodular in a and c:*

$$\frac{\partial^2 \ln \pi(a,n,c)}{\partial a \partial c} > 0 \quad \textit{and} \quad \frac{\partial^2 \ln S(a,n,c)}{\partial a \partial c} > 0.$$

Thus, every selection from $a_\pi(n,c)$ or $a_S(n,c)$ is increasing in c.

[7] Proposition 5 requires $\left.\frac{\partial S(a,n,c)}{\partial a}\right|_{a=\underline{a}} \neq 0$ and $\left.\frac{\partial S(a,n,c)}{\partial a}\right|_{a=\overline{a}} \neq 0$, which is generically true.

Propositions 6 and 7 imply that firm and consumer preferences are more likely to be aligned with either *few* or *many* firms in the market and with either *low* or *high* costs of production, as captured by the next corollary.

Corollary 8. *Assume* $\mu''(a) \geq 0$*. Then there exist* $\underline{n}, \overline{n} \in \mathbb{N} \cup \infty$ *and* $\underline{c}, \overline{c} \in \mathbb{R}_+ \cup \infty$ *such that firms' and consumers' most preferred* $a \in A$ *coincides for all* ($n \leq \underline{n}$ *or* $n \geq \overline{n}$) *and* ($c \leq \underline{c}$ *or* $c \geq \overline{c}$).

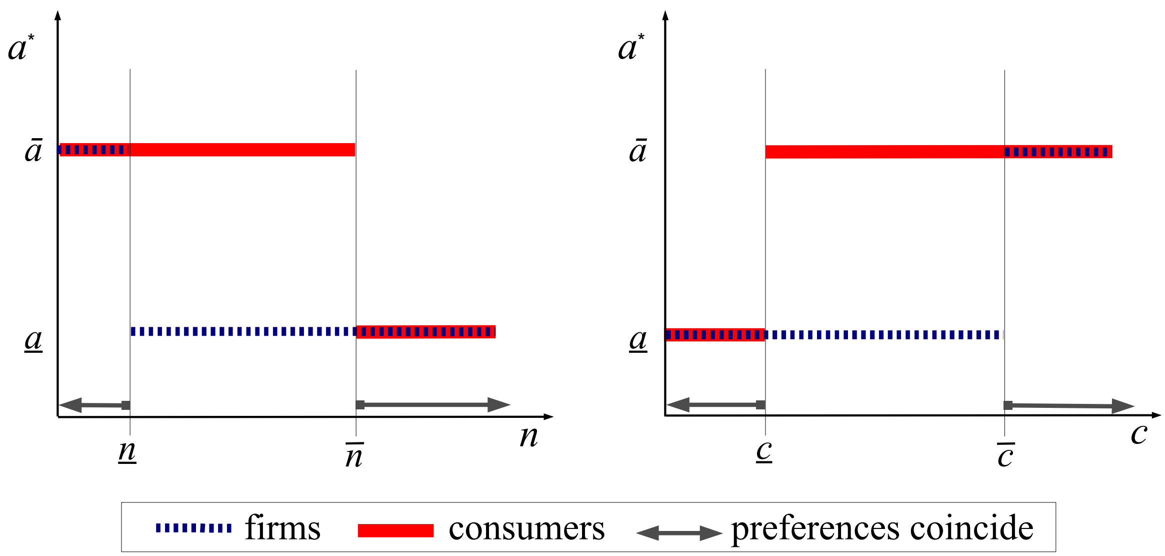

FIGURE 3.3. Preferred levels of advertising as a function of the number of firms and production costs

Figure 3.3 graphically illustrates the misalignment and potential harmony in advertising preferences shown in Propositions 4 through 7 and Corollary 8. First, both graphs exhibit Proposition 5, where firms' and consumers' preferences for advertising, indicated by the dashed and bold lines, respectively, always lie at the extreme levels of advertising, $\underline{a}$ and $\overline{a}$. However, by Proposition 4, consumers prefer higher levels of advertising than firms, indicated by the consumers' bolded preference line that always is higher (up to the maximum, $\overline{a}$) than the firms' dashed preference line in both graphs. In either graph, only the overlap of the two preference lines indicate alignment of preferences for advertising. In both graphs, we notice that there may exist an entry level of firms in the market and level of costs for which firms' and consumers' preferences do not align, indicated by separated dashed and bold preference lines. For example, one might imagine the two graphs superimposed, where there exists some level $n \in [\underline{n}, \overline{n}]$ and level $c \in [\underline{c}, \overline{c}]$ such that the firms' and consumers' preference lines do not overlap, i.e. advertising preferences do not align. This illustrates the main result of our analysis.

However, also notice in the graphs that for certain cost levels and number of firms in the market, advertising preferences do align. The left graph in Figure 3.3 displays Proposition 6. As the number of firms in the market decreases, firms' and consumers' preferences for high levels of advertising begin to align. Similarly, as the number of firms in the market increases, firms' and consumers' optimal levels of advertising decrease, ultimately aligning to the minimal level of advertising. Furthermore, we observe the first part of Corollary 8, where there exists an entry level into the market for which preferences stop aligning ($\underline{n}$) or begin to align ($\overline{n}$). Recall, the only general restriction on the entry level is $n \geq 2$.

The graph to the right in Figure 3.3 illustrates Proposition 7. As costs decrease, firms' and consumers' preferences for advertising decrease to the minimal level. Similarly, as costs increase, preferences align to the maximal level of advertising. As in the case of entry levels, there exists some cost levels for which preferences stop aligning ($\underline{c}$) and begin to align ($\overline{c}$), as given in Corollary 8.

4. Application to pharmaceutical advertising

We apply our results from above to DTCA in the pharmaceutical industry and the particular forms of regulated broadcast DTCA discussed in Section 2.1. Three major assumptions characterize our model. First, we assume advertising, the information dispersion parameter, is exogenously determined. Second, we assume goods are homogeneous. Third, we specify a constant elasticity of demand. Any loss in generality does not limit the application of our results to DTCA in the pharmaceutical industry.

First, our assumption of an exogenously determined advertising parameter complements the literature on co-opetition (cooperate-then-compete) (e.g. Brandenburg and Nalebuff (1996)). In our setting, the level of product information provided to consumers is agreed upon by firms before they compete in quantities, *à la* Cournot.[8] In the context of pharmaceutical drugs, we assume an industry consisting of the incumbent brand and generic entrants jointly choose a level of information dispersion that is realized through DTCA prior to competition.

For example, the pharmaceutical industry might explicitly cooperate on hype advertisements or agree upon a method to present positive and/or negative side effects. In recent years, aspirin suppliers began highlighting aspirin's benefits to the heart. In at least one instance, the FDA accused one of these aspirin suppliers (Bayer) of violating regulations in its advertisements of the health benefits associated with aspirin.[9]

While both brand name and generic firms advertise over-the-counter drugs, most advertising for prescription drugs is rendered by the brand name, with little participation by generic entrants (see Scherer (2000)). However, one might consider the

[8]For example, producers of milk compete outside their cooperation in the *Got Milk?* advertising campaign. Other trade associations similarly cooperate in advertising.

[9]See Warning Letter to Bayer, U.S. Food and Drug Administration (2008b).

selection of information disclosure as an industry decision, regardless of which firm or firms ultimately choose that level of information provision. In our model, we assume all firms advertise at the agreed upon level of advertising. If a low level of dispersion is optimal, a trade association or group of firms in the market might use DTCA to market a treatment though hype product information. Alternatively, to downplay the risks and negative effects of a drug, firms may tacitly or explicitly agree upon the real information they will disclose.

Next, with respect to our assumption of homogeneous goods, upon entry of generic firms into the pharmaceutical market, entrants must show bioequivalence to the patented drug (Reiffen and Ward (2005) and Frank and Salkever (1992)). Also, for a given drug, the quantity and strength of a dosage generally is consistent among suppliers (Meyerhoefer and Zuvekas (2008) and Sroka (2002)). Therefore, given the required bioequivalence of the compound and equivalent dosage of a drug supplied within a certain class, we assume a given drug produced in the market is homogeneous from the perspective of sellers and buyers.

Third, empirical analyses of the pharmaceutical industry frequently assume a constant elasticity of demand specification (e.g., Rizzo (1999), Caves, Whinston, Hurwitz, Pakes, and Temin (1991), and Pecorino (2002)). As may be the case with demand for a pharmaceutical drug, when observed prices and quantities do not vary widely, a constant elasticity demand specification is justified.

In Proposition 5, we extend Johnson and Myatt's result that firm profits are U-shaped in the information dispersion parameter and show that the same holds for consumers. In the context of pharmaceutical DTCA, this implies firms and consumers prefer advertisements that provide either the maximal level of information (real) or the minimum level of information (hype).

With respect to the maximal level of information, regulation pushing firms towards a niche market strategy by requiring the provision of more real information will result in an increase in the equilibrium price. Advertising increases the dispersion of consumer valuations, so firms optimally respond to the resulting increase in valuations of the remaining niche market participants by choosing a higher price. Our results in Proposition 4 highlight, however, that for a given number of firms in the industry, if firms chase a niche market by choosing higher levels of information, then consumers also prefer higher levels of information. Therefore, a niche market strategy, even if it results in higher equilibrium prices, might be preferred by consumers to an advertising strategy yielding lower equilibrium prices. While firms and consumers align in their preferences for extreme, rather than intermediate, levels of advertising, we justify the need for DTCA regulation by showing that firms and consumers do not necessarily prefer the *same* extreme level of advertising.

Proposition 4 also implies consumers always prefer more information than firms prefer. Therefore, under competition, firms might exogenously choose a minimal level of information, while consumers are better off with more real information. DTCA offers an example of how this misalignment of information preference can occur. Under one form of broadcast DTCA regulated by the FDA, firms are allowed to provide the drug name and intended use, but the FDA requires the advertisements provide ample *real* information (negative effects and risks) to consumers. As with the example of the Yaz® advertisements or recent campaigns related to the health benefits of aspirin, firms push the limits of regulation by downplaying negative effects and risks while hyping side benefits of a drug. Firms exhibit a preference for lower levels of information, while consumers might prefer more real information.

Numerous warning letters issued by the FDA to firms reveal the FDA's diligence in ensuring that firms provide more real information to consumers, per its DTCA requirements.[10] Yet, some forms of regulation of broadcast DTCA legally allow firms to pursue a mass market strategy through the provision of hype information under minimal restrictions. The FDA allows firms to hype a medical condition or name of the drug through "reminder advertisements" and "help-seeking advertisements." In this case, consumers may be better off with the provision of real information, and FDA regulation of these two forms of broadcast advertising may fail to protect consumers. Now, for prescription drugs, the end-user cannot access the pharmaceutical without a doctor's approval, in which case an argument can be made that the consumer could obtain more real information from the doctor.[11] However, in the case of advertising and labeling of over-the-counter drugs allowance of hype information may not be in the best interest of consumers.

While our results justify the need for regulating pharmaceutical advertising, we also show in Proposition 6 that for some level of entry into the industry, firm and consumer preferences for advertising align. As competition increases, consumers benefit more from a decrease in price than an increase in information. Therefore, as the number of generic firms entering the industry increases, firms prefer a mass market advertising strategy, and consumers agree with this provision of less product information about pharmaceuticals.

Similarly, in Proposition 7, as marginal cost decreases, firms and consumers align to a preference for low levels of advertising. In the market for pharmaceutical drugs,

[10] These letters are supplied by the Center for Drug Evaluation and Research Freedom of Information Office, published on the FDA's website.

[11] Yet, an argument also could be made that unless physicians educate themselves or participate in academic detailing (a process by which a third party administrator educates physicians about pharmaceuticals), then due to lack of information, the physician might not provide as much real information as consumers desire (see, e.g. Hill, Bunn, and Hawkins (2002)).

whether prescription or over-the-counter, if marginal costs of production are low, firms might prefer a mass market strategy. However, even with an aligning of preferences for low advertising, as marginal cost decreases, Figure 3.3 reveals that until a lower bound $\underline{c}$ is reached, consumers still prefer high levels of advertising. By showing bounds on entry into the market and marginal costs, we learn there exists a range in which consumer and firm preference for advertising do not align. For any level of advertising within these ranges, one might argue that regulation of advertising might push firms towards levels of information provision more closely desired by consumers.

5. Conclusion

We build upon the demand rotation framework of Johnson and Myatt (2006) by comparing firm and consumer preferences for advertising when the latter increases the inverse demand elasticity. Our goal is to investigate the alignment of firm and consumer preferences for advertising when firms *ex ante* agree upon a level of information provision, then compete in quantities in a market for a homogeneous good. We apply our results to advertising of pharmaceuticals, either through prescription drug DTCA or over-the-counter drug advertising.

We learn that under the demand rotation framework, not only do firms prefer the extreme levels of advertising, but so do consumers. However, their preferences for which extreme do no necessarily align. In particular, since we show consumers always prefer more advertising than firms, then firms may optimally provide minimal levels of advertising, while consumers prefer maximal levels of advertising. With regards to advertisements for prescription or over-the-counter drugs, advertising that allows firms to provide only hype information may not result in the levels of advertising preferred by consumers, justifying greater or stricter advertising regulation. However, despite this possible misalignment of preferences, we also show that either few or many firms in the market and either cheap or expensive technologies can alleviate these differences without regulation. Therefore, our analysis reveals that firm advertising behavior might not serve consumers' best interests, unless certain market attributes hold.

6. Appendix C: Ambiguous results of a general analysis

For equilibrium total quantity $Z(a,\cdot)$, equilibrium profit for firm i is:

$$\pi_i(Z(a,\cdot)) = \frac{Z(a,\cdot)}{n}\left(P(Z(a,\cdot)) - c\right).$$

In equilibrium,

$$c = \frac{\partial P(Z(a,\cdot)}{\partial Z}\frac{Z(a,\cdot)}{n} + P(Z(a,\cdot)).$$

One can easily show:

$$\frac{d\pi_i(Z(a,\cdot))}{da} = \frac{Z(a,\cdot)}{n}\left[\frac{dP(Z(a,\cdot)}{da} - \frac{\partial P(Z(a,\cdot)}{\partial Z}\frac{dZ(a,\cdot)}{da}\frac{1}{n}\right].$$

Without knowing the effect of advertising on industry quantity, one cannot sign the effect of advertising on per-firm profit under oligopoly competition. The reaction of equilibrium total quantity to an increase in advertising will depend on how the slope of the inverse demand curve reacts to an increase in advertising. However, a rotation ordering places no restriction on how the slope of demand responds to advertising.

Therefore, since rotation ordering is independent of a general restriction on slope, Johnson and Myatt's monopoly analysis does not generally extend to Cournot oligopoly, due to these strategic effects. We must place restrictions on the slope of demand. For this reason, we assume consumer valuations are distributed such that the index of distribution functions, a, also is the measure of consumer sensitivity to price changes.

7. Appendix D: Chapter 3 proofs

Proof of Proposition 2: Johnson and Myatt (2006) provide an expression for aggregate output at the symmetric equilibrium. The fact that no asymmetric equilibrium exists follows by Amir and Lambson (2000). The expression for $P\left(a,n,c\right)$ follows by substituting $Z\left(a,n,c\right)$ in $P_a\left(Z\right)$.□

Proof of Proposition 3: Johnson and Myatt (2006) provide an expression for Cournot profits at the unique and symmetric equilibrium. Consumer surplus is defined as $S\left(a,n,c\right) := \int_0^{Z(a,n,c)} P_a\left(Z\right)dZ - P_a\left(Z\right)Z\left(a,n,c\right)$. The result follows by substituting $P_a\left(Z\right)$ and $Z\left(a,n,c\right)$ in the latter expression.□

Proof of Proposition 4: It follows from Proposition 3 that

$$S\left(a,n,c\right) = \left[n^2/\left(1-a\right)\right]\pi\left(a,n,c\right).$$

We next show that for all $a, a' \in A$ with $a > a'$,

$$\pi\left(a,n,c\right) - \pi\left(a',n,c\right) \geq 0 \Longrightarrow S\left(a,n,c\right) - S\left(a',n,c\right) > 0. \tag{3.1}$$

The right hand side of (3.1) is equivalent to

$$\pi\left(a,n,c\right) \geq \pi\left(a',n,c\right). \tag{3.2}$$

Note $\left[n^2/\left(1-a\right)\right] > \left[n^2/\left(1-a'\right)\right]$. Since $\pi\left(a,n,c\right) > 0$ and $\pi\left(a',n,c\right) > 0$, then

$$\left[n^2/\left(1-a\right)\right]\pi\left(a,n,c\right) > \left[n^2/\left(1-a'\right)\right]\pi\left(a',n,c\right). \tag{3.3}$$

which is the left hand side of (3.1). This proves the claim.

The fact that any element in $a_S(n,c)$ is larger than every element in $a_\pi(n,c)$ follows by Amir (2005), Theorem 12.□

Proof of Proposition 5: Johnson and Myatt (2006) show if $\mu''(a) \geq 0$, then $\pi(a,n,c)$ is convex in a. The fact it is strictly convex follows directly from their proof. Now

$$\left.\frac{\partial^2 S(a,n,c)}{\partial a^2}\right|_{\frac{\partial S(a,n)}{\partial a}=0} = n^2\left[\frac{2}{(1-a)^3}+\frac{a}{(1-a)^2}\right]\pi(a,n,c)+\left[\frac{n^2}{(1-a)}\right]\frac{\partial^2 \pi(a,n,c)}{\partial a^2}.$$

Since $\mu''(a) \geq 0$, then quasiconvexity of $S(a,n,c)$ in a follows by the convexity of $\pi(a,n,c)$. The second claim follows directly from the last two.□

Proof of Proposition 6: Computing the cross partial derivatives of $\ln\pi(a,n,c)$ and $\ln S(a,n,c)$ with respect to a and n we get

$$\frac{\partial^2 \ln\pi(a,n,c)}{\partial a \partial n} = \frac{\partial^2 \ln S(a,n,c)}{\partial a \partial n} = -\frac{n-1}{n(n-a)^2} < 0.$$

The fact that any selection from $a_\pi(n)$ and $a_S(n)$ is decreasing in n follows by Amir (2005), Theorem 12.□

Proof of Proposition 7: The proof is similar to that of Proposition 6, so we omit it.□

Proof of Corollary 8: If $\mu''(a) \geq 0$, $a_\pi(n,c) \in \{\underline{a},\overline{a}\}$ and $a_S(n,c) \in \{\underline{a},\overline{a}\}$ by Proposition 5. By Proposition 4 every element in $a_S(n,c)$ is higher than any element in $a_\pi(n,c)$, thus if consumers prefer $\underline{a}$ over $\overline{a}$ the same holds for the firms. Define $\overline{n}$ as the smallest integer that satisfies $S(\underline{a},\overline{n},c) \geq S(\overline{a},\overline{n},c)$. By Proposition

6 every selection from $a_S(n,c)$ decreases in n. Thus if $a_S(\overline{n},c)=\underline{a}$ it must hold that $a_S(n,c)=\underline{a}$ and $a_\pi(\overline{n},c)=\underline{a}$ for all $n\geq\overline{n}$.

It also true that if firms prefer $\overline{a}$ over $\underline{a}$ then the same holds for consumers. Define $\underline{n}$ as the largest integer that satisfies $\pi(\overline{a},\underline{n},c)\geq\pi(\underline{a},\underline{n},c)$. By Proposition 6 every selection from $a_S(n,c)$ decreases in n. Thus if $a_\pi(\underline{n},c)=\overline{a}$ it must hold that $a_\pi(n,c)=\overline{a}$ and $a_\pi(n,c)=\overline{a}$ for all $n\leq\underline{n}$.

The proof for costs is similar, so we omit it.□

References

Ackerberg, D. (2001): "Empirically Distinguishing Informative and Prestige Effects of Advertising," *RAND Journal of Economics*, 32(2), 316–333.

Aislabie, C. J., and C. A. Tisdell (1988): "Profit Maximization and Marketing Strategies: Demand Rotation and Social Influences," *Managerial and Decision Economics*, 9(1), 77–82.

Amir, R. (2005): "Supermodularity and Complementarity in Economics: An Elementary Survey," *Southern Economic Journal*, 71, 636–660.

Amir, R., and V. Lambson (2000): "On the Effects of Entry in Cournot Markets," *Review of Economic Studies*, 67, 235–254.

Amir, R., and N. Lazzati (2010): "Network Effects, Market Structure and Industry Performance," mimeo.

Apollo Group Inc v. Policemen's Annuity and Benefit Fund of Chicago (2011): 131 S.Ct. 1602, 79 USLW 3310, 79 USLW 3511, 79 USLW 3512 (U.S.2011).

Bagwell, K. (2007): *Handbook of Industrial Organization* vol. 3, chap. 28: The Economic Analysis of Advertising, pp. 1701–1844.

Basic v. Levinson (1988): 485 U.S. 224,108 S. Ct. 978,99 L. Ed. 2d 194,1988 U.S.

Bayer Pharmaceuticals (Creator), Unquid (Poster) (2009, February 26): "Yaz FDA required "clear up"," [Video] retrieved from Youtube.com: http://www.youtube.com/watch?v=EO-G8O0lHq0, (last accessed April 2011).

Bhagat, S., and R. Romano (2002): "Event Studies and the Law: Part II: Empirical Studies of Corporate Law," *American Law and Economics Review*, 4(2), 380–423.

Brandenburg, A., and B. Nalebuff (1996): *Co-opetition*. Doubleday, New York.

Brav, A. (2000): "Inference in Long-Horizon Event Studies: A Bayesian Approach with Application to Initial Public Offerings," *The Journal of Finance*, 55(5), 1979–2016.

Brown, S. J., and J. B. Warner (1980): "Measuring Security Price Performance," *Journal of Financial Economics*, 8, 205–258.

——— (1985): "Using daily stock returns : The case of event studies," *Journal of Financial Economics*, 14(1), 3–31.

Buchanan, J. M., and Y. J. Yoon (2000): "Symmetric Tragedies; Commons and Anticommons," *Journal of Law and Economics*, 43(1).

Carmichael, L., and W. B. MacLeod (2003): "Caring About Sunk Costs: A Behavioral Solution to Holdup Problems with Small Stakes," *Journal of Law, Economics, and Organization*, 19(1), 106–118.

Caves, R. E., M. D. Whinston, M. A. Hurwitz, A. Pakes, and P. Temin (1991): "Patent Expiration, Entry, and Competition in the U.S. Pharmaceutical Industry," *Brookings Papers on Economic Activity. Microeconomics*, pp. 1–66.

Chin, M.-C., and R. Weaver (2002): "Does Generic Advertising Wrap Demand Curvature?," Selected Paper for AAEA Annual Meeting, Long Beach, California July 28-31, 2002.

Cichello, M., and D. J. Lamdin (2006): "Event Studies and the Analysis of Antitrust," *International Journal of the Economics of Business*, 13(2), 229–245.

Cornell, B., and R. G. Morgan (1990): "Using Finance Theory to Measure Damages in Fraud on the Market Cases," *UCLA Law Review*, 37, 883–924.

Cournot, A. (1838): *Recherches sur les Principes Mathematizues de la Theorie des Richesses*. L. Hachette. English translation by N.T. Bacon in Economic Classics, New York, NY: Macmillan, 1897; reprinted by Augustus M. Kelly, 1960.

Cowan, S. (2004): "Demand Shifts and Imperfect Competition," University of Oxford Department of Economics Discussion Paper Series, No. 188.

Dari-Mattiacci, G., and F. Parisi (2006): "Substituting Complements," *Journal of Competition Law and Economics*, 2(3), 333–347.

Daubert v. Merrell Dow Pharmaceuticals, Inc. (1993): (92-102), 509 U.S. 579, 113 S.Ct. 2786 (U.S.Cal.,1993).

Dixit, A., and V. Norman (1978): "Advertising and Welfare," *The Bell Journal of Economics*, 9(1), 1–17.

Dorfman, R., and P. O. Steiner (1954): "Optimal Advertising and Optimal Quality," *The American Economic Review*, 44(5), 826–836.

Dyl, E. A. (1999): "Estimating Economic Damages in Class Action Securities Fraud Litigation," *Journal of Forensic Economics*, 12(1), 1–11.

Edgeworth, F. Y. (1925): "The Pure Theory of Monopoly," in Papers Relating to Political Economy. for the Royal Econ. Soc. Vol 1. F. Y. Edgeworth (eds). London: Macmillan Co.

Ellet, C. J. (1839): *An Essay on the Laws of Trade In Reference to the Works of Internal Improvement in the United States*. P.D. Bernard, Richmond, Virginia.

Ellickson, R. C. (1993): "Property in Land," *Yale Law Journal*, 102, 1315–1400.

Ellingsen, T., and M. Johannesson (2004): "Is There a Hold-up Problem?," *Scandinavian Journal of Economics*, 106, 475–494.

Fama, E. F., L. Fisher, M. C. Jensen, and R. Roll (1969): "The Adjustment of Stock Prices to New Information," *International Economic Review*, 10(1), 1–21.

Feinberg, Y., and M. I. Kamien (2001): "Highway robbery: complementary monopoly and the hold-up problem," *International Journal of Industrial Organization*, 19(10), 1603–1621.

Fener v. Belo Corp. (2008): 560 F.Supp.2d 502 (N.D.Tex.,2008).

Fennell, L. A. (2010): *Research Handbook on the Economics of Property Law* chap. 2: Commons, Anticommons, Semicommons. Kenneth Ayotte, Henry E. Smith (eds). Edward Elgar.

Frank, R. G., and D. S. Salkever (1992): "Pricing, Patent Loss and the Market for Pharmaceuticals," *Southern Economic Journal*, 59(2), 165–180.

Gardner, R., N. Gaston, and R. T. Masson (August 2002): "Tolling the Rhine in 1254: Complementary Monopoly Revisited," *Unpublished paper*.

Gelbach, J., E. Helland, and J. Klick (2010): "Valid Inference in Single-Firm, Single-Event Studies," Robert Day School Working Paper No. 2009-17.

Gelfand, A. E., and J. M. Sfiridis (1996): *Advances in Econometrics* vol. II(A), chap. 11: Bayesian Analysis of Financial Event Studies Data, pp. 25–62. JAI Press, Inc.

Geweke, J. (2005): *Contemporary Bayesian Econometrics and Statistics*. John Wiley & Sons, Inc., New Jersey.

Graves, P., and R. Sexton (2006): "Demand and Supply Curves: Rotations versus Shifts," *Atlantic Economic Journal*, 34, 361–364.

Gul, F. (2001): "Unobservable Investment and the Hold-Up Problem," *Econometrica*, 69, 343376.

Halmenschlager, C., A. Mantovani, and M. Troege (2011): "Demand Expansion and Elasticity Improvement as Complementary Marketing Goals," *The Manchester School*, 79, 145–158.

Hamilton, S. F. (1999): "Demand Shifts and Market Structure in Free-Entry Oligopoly Equilibria," *International Journal of Industrial Organization*, 17(2), 259–275.

Hawkins, J. R., and N. Lazzati (2011): "Misalignment of Firm and Consumer Preferences for Advertising," Working Paper.

Heller, M. (1998): "The Tragedy of the Anticommons: Property in the Transition from Marx to Markets," *Harvard Law Review*, 111, 621–688.

Heller, M. (2008): *The Gridlock Economy: How Too Much Ownership Wrecks Markets, Stops Innovation, and Costs Lives*. Basic Books.

Heller, M. A., and R. S. Eisenberg (1998): "Can Patents Deter Innovation? The Anticommons in Biomedical Research," *Science*, 280(5364), 698–701.

Hill, C. D., D. N. Bunn, and J. R. Hawkins (2002): "Stretching the Managed Care Dollar in the New Millennium: The Practice of Detailing Primary Care Physicians," *Managed Care Quarterly*, 10(2).

Iizuka, T. (2004): "What Explains the Use of Direct-to-Consumer Advertising of Prescription Drugs?," *The Journal of Industrial Economics*, 52(3), 349–379.

Ikenberry, D., J. Lakonishok, and T. Vermaelen (1995): "Market Underreaction to Open Market Share Repurchases," *Journal of Financial Economics*, 39, 181–208.

Illinois Attorney General (2009): "Madigan Reaches Settlement With Bayer for Violating 2007 Agreement to Cease Deceptive Marketing Practices," http://illinoisattorneygeneral.gov/pressroom/.

In re Bare Escentuals, Inc. Securities Litigation (2010): 745 F.Supp.2d 1052 (N.D.Cal.,2010).

In re Imperial Credit Industries, Inc. (2008): WL 4346785 (C.D.Cal.,2008).

In re Omnicom Group, Inc. Securities Litigation (2010): 597 F.3d 501 (C.A.2 (N.Y.),2010).

In re Williams Securities Litigation-WCG Subclass (2009): 558 F.3d 1144 (C.A.10 (Okla.),2009).

Jeffreys, H. (1939): *Theory of Probability*, International series of monographs on physics. Clarendon Press, reprinted 1998, 3rd edition.

Johnson, J., and D. Myatt (2006): "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 96, 756–784.

Kass, R. E., and A. E. Raftery (1995): "Bayes Factors," *Journal of the American Statistical Association*, 90(430), pp. 773–795.

Kinnucan, H., and Y. Zheng (2004): "Advertising's Effect on the Market Demand Elasticity: A Note," *Agribusiness*, 20(2), 181–188.

Klick, J., and R. H. Sitkoff (2008): "Agency Costs, Charitable Trusts, and Corporate Control: Evidence from Hershey's Kiss-Off," *Columbia Law Review*, 108(4).

Kumho Tire Co. v. Carmichael (1999): (97-1709) 526 U.S. 137.

Lewis, T. R., and D. E. M. Sappington (1991): "All-or-Nothing Information Control," *Economics Letters*, 37(2), 111–113.

Ling, D. C., E. R. Berndt, and M. K. Kyle (2002): "Deregulating Direct-to-Consumer Marketing of Prescription Drugs: Effects on Prescription and Over-the-Counter Product Sales," *Journal of Law and Economics*, 45(S2), 691–723.

Llanes, G., and S. Trento (2009): "Anticommons and Optimal Patent Policy in a Model of Sequential Innovation," *Harvard Business School Working Papers 09-148, Harvard Business School.*

Lueck, D., and T. J. Miceli (2007): *Handbook of Law and Economics - Vol 2* chap. 3: Property Law. Michael Polinsky and Steven Shavell (eds). North-Holland.

Machlup, F., and M. Taber (1960): "Bilateral Monopoly, Successive Monopoly, and Vertical Integration," *Economica*, 27(106), 101–119.

Meyerhoefer, C., and S. Zuvekas (2008): "The Shape of Demand: What Does It Tell Us about Direct-to-Consumer Marketing of Antidepressants?," *The B.E. Journal of Economic Analysis & Policy*, 8(2).

Michelman, F. I. (1982): *Ethics, Economics and the Law.* J. Roland Pennock and John W. Chapman (eds). New York University Press.

Nichols, L. M. (1985): "Advertising and Economic Welfare," *The American Economic Review*, 75(1), 213–218.

NIHCM Foundation (2001): "Prescription Drugs and Mass Media Advertising, 2000," Report.

Okuguchi, K., and F. Szidarovszky (2005): “Changes in Demand Function in Cournot Oligopoly,” *Pacific Economic Review*, 10(3), 371–378.

Oscar Private Equity Investments v. Allegiance Telecom, Inc. (2007): 487 F.3d 261, 269 (5th Cir. 2007).

Palumbo, Francis B.; Mullins, C. D. (2002): “The Development of Direct-to-Consumer Prescription Drug Advertising Regulation,” *Food & Drug Law Journal*, 57.

Parisi, F., and B. Depoorter (2003): *The Economics of Copyright: Developments in Research and Analysis* chap. The Market for Intellectual Property: The Case of Complementary Oligopoly, p. 261. W. Gordon, R. Watt (eds). Edward Elgar.

Parisi, F., B. Depoorter, and N. Schulz (2005): “Duality in Property: Commons and Anticommons,” *International Review of Law and Economics*, 25(4).

Parisi, F., N. Schulz, and B. Depoorter (2004): “Simultaneous and Sequential Anticommons,” *European Journal of Law and Economics*, 17(2).

Pecorino, P. (2002): “Should the US Allow Prescription Drug Reimports from Canada?,” *Journal of Health Economics*, 21(4), 699–708.

Pines, W. L. (1999): “A History and Perspective on Direct-to-Consumer Promotion,” *Food and Drug Law Journal*, 54(4), 489–518.

Quirmbach, H. C. (1988): “Comparative Statics for Oligopoly: Demand Shift Effects,” *International Economic Review*, 29(3), 451–459.

Reiffen, D., and M. R. Ward (2005): “Generic Drug Industry Dynamics,” *The Review of Economics and Statistics*, 87(1), 37–49.

Rizzo, J. A. (1999): “Advertising and Competition in the Ethical Pharmaceutical Industry: The Case of Antihypertensive Drugs,” *Journal of Law and Economics*, 42(1), 89–116.

Salinger, M. A. (1989): “The Meaning of ”Upstream” and ”Downstream” and the Implications for Modeling Vertical Mergers,” *The Journal of Industrial Economics*, 37(4), 373–387.

Scherer, F. (2000): *Handbook of Health Economics*chap. 25: The Pharmaceutical Industry. Cuyler, A.J., Newhouse, J.P. (eds). Elsevier, North Holland, Amsterdam.

Schultz, D. (2008): “FDA Oversight of Direct-to-Consumer Advertising of Medical Devices,” Statement before the Senate Special Committee on Aging, September 17, 2008.

Sfiridis, J., and A. Gelfand (2002): “A Survey of Sampling Based Bayesian Analysis of Financial Data,” *Applied Mathematical Finance*, 9, 273–291.

Shapiro, C. (1980): “Advertising and Welfare: Comment,” *The Bell Journal of Economics*, 11(2), 749–752.

Singer, N. (2009): “A Birth Control Pill That Promised Too Much,” *New York Times*.

Sonnenschein, H. (1968): “The Dual of Duopoly Is Complementary Monopoly: or, Two of Cournot’s Theories Are One,” *The Journal of Political Economy*, 76(2), 316–318.

Spengler, J. J. (1950): “Vertical Integration and Antitrust Policy,” *Journal of Political Economy*, 58, 347.

Sroka, C. J. (2002): *The Pharmaceutical Industry*chap. 2: Prescription Drugs: Pricing Differences between Insured and Uninsured Consumers, p. 32. Ethan N. Parvis (ed). Nova Science Publishers, Inc.

Stackelberg, H. (1934): *Marktform und Gleichgewicht*. Springer-Verlag. English translation, entitled The Theory of the Market Economy, published 1952, Oxford University Press.

Stanford Law School Securities Class Action Clearinghouse (2011): http://securities.stanford.edu(last accessed 04/2011).

Steen, F., and K. G. Salvanes (1999): “Testing for market power using a dynamic oligopoly model,” *International Journal of Industrial Organization*, 17(2), 147–177.

Sunpower Corp. Systems v. Sunlink Corp. (2009): WL 1657987 (N.D.Cal.,2009).

Sussman, A. (2008): “Damages and Loss Causation NEW YORK CITY BAR Securities Litigation, During the Credit Crisis: Current Developments and Strategies,” Presentation.

Tirole, J. (1988): *The Theory of Industrial Organization*. The MIT Press, Cambridge, Massachusetts and London, England.

U.S. Code 21 C.F.R. §202 (1997, August): *FDA, Guidance for Industry: Consumer-directed Broadcast Advertisements* (Washington, D.C.: FDA).

U.S. Food and Drug Administration (2008a): “Warning Letter to Bayer Healthcare Pharmacueticals, Inc.,” http://www.fda.gov, (last accessed April 2011).

——— (2008b): "Warning Letter to Bayer Healthcare Pharmacueticals, Inc.," http://www.fda.gov/ICECI/EnforcementActions/WarningLetters, (last accessed April 2011).

Vanneste, S., A. Van Hiel, F. Parisi, and B. Depoorter (2006): "From 'Tragedy' to 'Disaster': Welfare Effects of Commons and Anticommons Dilemmas," *International Review of Law and Economics*, 26, 104–122.

Williamson, O. E. (1975): *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press, New York.

Zellner, A. (1975): "Bayesian Analysis of Regression Error Terms," *Journal of the American Statistical Association*, 70(349), 138–144.

Zheng, Y., H. Kinnucan, and H. Kaiser (2010): "Measuring and testing advertising-induced rotation in the demand curve," *Applied Economics*, 42(13), 1601–1614.