Three Essays on Strategic Behavior and Policy

by

Danial Asmat

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in the University of Michigan
2015

Doctoral Committee:

Professor Valerie Y. Suslow, Co-Chair
Adjunct Professor Margaret C. Levenstein, Co-Chair
Professor Daniel A. Ackerberg
Professor Francine Lafontaine
Assistant Professor Brian Wu

To my mother, Hina, and my father, Parvez.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF APPENDICES

CHAPTER 1

Collusion Along the Learning Curve: Theory and Evidence from the Semiconductor Industry

This paper posits a theory of collusion with learning-by-doing and multiproduct competition and tests it with an explicit price fixing cartel. It introduces a novel repeated game model in which firms collude across products with time-varying learning. The model shows that collusion is harder to sustain in the early stage of a product life cycle, when learning is high, than in the later stage of a life cycle, as learning declines. By raising the price of an older product generation, successful collusion will shift demand toward a frontier product generation. The model's predictions are tested using data from the Dynamic Random Access Memory (DRAM) industry, which features frequent, repeated learning curves with each new product generation. Major manufacturers in the DRAM market also recently pled guilty to charges of illegal price fixing. Empirical analysis exploits variation between collusive and competitive time periods to show that collusion increased price by up to 25% for products at the end of their life cycle. In contrast, it reduced the price of new generations by up to 70%. Firms increased their output of new generations while cutting their output of older generations. The evidence implies that the DRAM cartel successfully colluded on mature products, which shifted demand toward newer products. Increased output of new products in collusion relative to competition allowed firms to learn faster in the collusive period.

1.1 Introduction

Since the turn of the century, cartels discovered in multibillion dollar high-technology markets have resulted in several of the largest antitrust fines in U.S. history.[1] Global manufacturers of LCD panels, cathode ray tubes, optical disk drives, and three different types of memory chips have settled criminal or civil claims for price fixing. These markets feature several characteristics long known to enable collusion, among them product homogeneity, cooperative research and development, and high barriers to entry. Yet they also share two additional characteristics. First, manufacturing displays learning-by-doing: a firm's cumulative output reduces the marginal cost of its future output. Second, firms steadily release new product generations based on technological advancements, such as those predicted by Moore's Law. Multiple product generations therefore overlap at the same time on the market.[2]

The discovery of such cartels raises a natural question: what is the impact of learning-by-doing on the damage attributable to collusion? Collusion is well understood to generate inefficiency by raising price above marginal cost. At first glance, learning raises this inefficiency: by restricting output, firms forfeit some of the gains to learning and inhibit cost reduction. This paper's contribution is to highlight that collusive equilibria in high-technology markets are determined by the rate at which firms learn and the extent of multiproduct demand linkage. It develops a repeated game model of learning with multiproduct competition to evaluate collusive equilibria relative to competitive equilibria at the product general level. It then uses data from an illegal price fixing cartel in the Dynamic Random Access Memory (DRAM) market to find evidence consistent with both the theory and its mechanism.

The model builds three insights in succession. First, collusion is at least as effective

[1] AU Optronics, LG Display, and Samsung Electronics have each received penalties of $300 million or more since 2006. AU Optronics' $500 million fine is the largest Sherman Act corporate fine to date; see http://www.justice.gov/atr/public/criminal/sherman10.html.

[2] In 2014, for exampe, buyers could choose between 2Gb, 4Gb and other generatons of chips. See http://www.forbes.com/sites/jimhandy/2014/06/27/dram-asps-soften-is-that-important/.

in an older product generation as a newer generation. This is because firms learn more early in the life cycle than later, and learning increases the incentive to defect from a collusive equilibrium. Second, collusion in the older generation shifts demand to the newer generation, which is an (imperfect) substitute. Third, if the demand shift is great enough for the newer generation, firms sell *more* of its units during collusion than competition. Counterintuitively, learning then creates efficiency *gains* from collusion, and such gains are transferred to buyers through lower prices if firms pass on enough cost savings.

DRAM presents an ideal application for several reasons. The product's explosive innovation has helped fuel the computer and electronics revolution, making it a critical part of the world's economy. DRAM production is a classic example of learning-by-doing: output allows firms to reduce cost-per-chip by reducing the rate of defective chips in manufacturing. Moreover, chip generations are released every two to three years, and the learning process repeats with each new generation. Because the product life cycle is several years, multiple generations overlap at any given time.

These three features—collusion, learning, and multiproduct competition—allow the model to be naturally tested. I employ firm-level data on the DRAM market before, during, and after dates of admitted cartel activity.[3] I estimate the change in price from competition to collusion separately for five DRAM generations active during the cartel period. I identify the cartel overcharge by exploiting variation between cartel and non-cartel time periods as well as variation in the stage of the product life cycle during the cartel time period.

The empirical results are strikingly consistent with the model's predictions. In the two most mature cartel generations, 4Mb and 16Mb, the overcharge is positive and as high as 26%. The overcharge among newer 64Mb, 128Mb and 256Mb generations, however, is sharply negative and as low as -70%. Results are robust to industry-wide capacity utilization and heterogeneity in underlying DRAM technology. Output data implies that firms sold significantly more 128Mb and 256Mb chips during the cartel period than they would

[3] Dates are based on publicly available litigation evidence from the Department of Justice (DOJ), European Commission (EC), and Noll (2014). See Section 1.3 for a full description.

otherwise, consistent with a demand shift toward frontier generations and more overall learning. Overall, the cartel's welfare effect was ambiguous.

This paper bridges a gap between two distinct groups of research in industrial organization: learning-by-doing and collusion. It contributes to both theory and empirical evidence in each of these two areas. It extends the theory of oligopolistic competition with learning to include the possibility of tacit or explicit coordination. Similarly, it adds learning within a realistic multiproduct setting to standard models of repeated game collusion. On the empirical side, it extends literature on the learning-intensive semiconductor industry as it enters an era of increased concentration, intellectual property disputes and antitrust scrutiny. It also joins literature that studies a "hard core" price fixing cartel to isolate the impact of collusion on welfare, which is otherwise difficult.

The model of learning-by-doing is closest to Siebert (2010) and Fudenberg and Tirole (1983): firms choose quantities of one product in a learning period knowing that they gain a cost reduction as a function of output in a post-learning period.[4] Products generations are imperfect substitutes for one another. The model is embedded within an infinite horizon game, as in Cabral and Riordan (1994) and Besanko et al. (2010). It differs from existing models of learning-by-doing in focusing on collusive equilibria rather than long-run industry dynamics such as increasing dominance, natural monopoly and predatory pricing.[5]

The model's fundamental contribution is to illustrate how collusion on the product market with learning and multiproduct demand linkages can lead to welfare-enhancing as well as welfare-reducing outcomes. Brod and Shivakumar (1999) and Fershtman and Pakes (2000) find a range of such equilibria when firms invest in R&D. It also adds to literature that assesses the sustainability of collusion with a dynamic variable more generally (Benoit and Krishna, 1987, Davidson and Deneckere, 1990 and Compte, Jenny and Rey, 2002).

[4]The Cournot framework is apt because capacity for semiconductors and other high-technology products is fixed within a product life cycle. Other quantity-setting models of learning include Spence (1981), Ghemawat and Spence (1985), Dasgupta and Stiglitz (1988) and Cabral and Riordan (1997).

[5]Mookherjee and Ray (1991) propose a model of collusion with scale economies and learning-by-doing. I expand on their results by adding demand linkages between product generations in a quantity-setting framework, which characterize numerous high-technology industries with capacity constraints.

These papers show that adding dynamics to standard models of collusion can fundamentally change equilibrium incentives and welfare outcomes.

The empirical analysis adds to a number of studies that test the theory of oligopolistic competition with learning using empirical data from the semiconductor industry. Irwin and Klenow (1994) specify a quantity-setting model to estimate that a firm's marginal cost of chip production is reduced by an average of 20% following a 100% increase in output, but that there is no significant learning effect between new product generations. Zulehner (2003) and Siebert and Zulehner (2013) examine the rate of learning and competitiveness in DRAM using data from the 1970s to the mid-1990s, before firms began explicitly colluding.[6] Gardete (2014) uses more recent data to estimate a structural model of imperfect information and finds that DRAM manufacturers share demand information in equilibrium.[7]

Finally, this paper contributes to empirical evidence of how cartels meet collusive incentive compatibility constraints (Levenstein, 1997; Scott Morton, 1997; Genesove and Mullin, 2001; Roller and Steen, 2006; Mariuzzo and Walsh, 2013; and Clark and Houde, 2013, 2014) by explicitly accounting for learning and multiproduct competition.[8] Unlike other studies that directly estimate cartel damages (Porter and Zona, 1999; Bolotova, Connor and Miller, 2008), it finds strong evidence that the DRAM cartel benefited some consumers through lower prices while harming others through higher prices.[9] It remains vital to understand the welfare impacts of collusion using data from modern cartels, which operate in vastly different legal and technological settings from those in the past.

The remainder of the paper is outlined as follows. Section 1.2 presents a multi-period,

[6]Several additional studies have found evidence consistent with significant learning effects in DRAM and other memory chips. See Dick (1991), Nye (1996), Gruber (1998), and Cabral and Leiblein (2001).

[7]The role of memory and microprocessor chips, as well as technology standards in memory chips, is explained in Section 3.2.1.

[8]See Levenstein and Suslow (2006) for a broader review of empirical literature examining incentive compatibility in cartels.

[9]Asker (2010) is a notable exception in estimating damages from a bidding ring for stamp collection auctions. He finds that bidders have incentives to inflate their private valuations in knockout auctions, sometimes resulting in higher stamp payments to sellers relative to the competitive counterfactual.

infinite horizon model of collusion over the product life cycle and generates testable predictions for prices during collusion. Section 1.3 and Section 1.4 describe the features of the industry and the data for estimation, respectively. Section 1.5 examines price and output effects of collusion between generations. I conclude in Section 1.6 by discussing the results and implications for antitrust policy.

1.2 Theoretical Model

I embed a discrete time quantity-setting model of learning based on Siebert (2010) and Fudenberg and Tirole (1983) into an infinitely-repeated duopoly game. Two firms, i and j, produce two products $k_r, r \in \{1, 2\}$. The game has three phases, and product k_{r+1} enters exogenously one phase after product k_r's entry. Firms compete in quantities and face linear demand for products that are imperfect substitutes.

Learning occurs between the game's first and second phases: output in the first phase reduces marginal cost in the second phase at the rate $\lambda \in [0, 1]$. Between the second and third phases, any further learning effects diffuse to both firms. From phase three onward, firms play a static Cournot game with identical costs. Marginal cost is c_1 for both firms in phase one, $c_{i2} = max\ \{c_1 - \lambda q_{ik1},\ c_3\}$ for firm i in phase two, and $c_3 \equiv \underline{c} < c_1$ in phase three to infinity. Profits are discounted at a common rate $\delta \in [0, 1]$ per period.

Firms simultaneously select quantities q_{ikt} at the start of each period, and the total quantity Q_{kt} determines the market-clearing price P_{kt}. After each period and before the next, firms observe the quantity decision of the other with perfect information.

This model captures the salient features of the product life cycle of DRAM and many other computer components.[10] Every several years, chip makers invest in costly new plants with fixed capacity.[11] Firms enter each DRAM product generation with high cost, reduce

[10]Examples include microprocessors, static random access memory (SRAM), flash memory, hard disk drives (HDD's), solid state drives (SSD's), optical disk drives (ODD's), servers, LCD panels, and graphics cards.

[11]Building a new plant to increase capacity requires about two years ("Memory Indus-

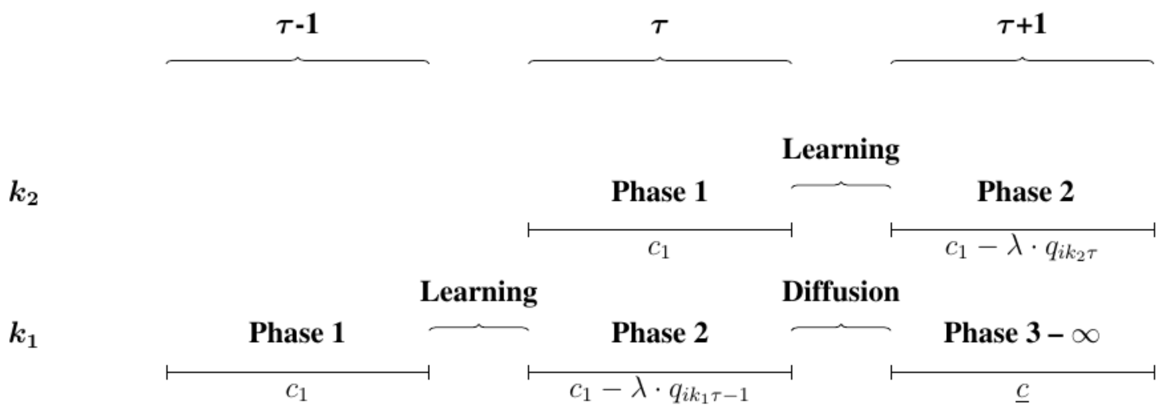

Figure 1.1: Timing of Product Life Cycles and Learning

cost by learning as a function of cumulative output, exploit a cost advantage over rivals if it develops, and eventually exhaust learning. This process repeats for every product generation.[12]

1.2.1 Model Preliminaries

Assume that demand for DRAM is linear and generations are imperfect substitutes:

$$P_{kt} = \begin{cases} a - \eta_2 Q_{k_2t} - \gamma_2 Q_{k_1t} & \text{if } k = k_2 \\ b - \eta_1 Q_{k_1t} - \gamma_1 Q_{k_2t} & \text{if } k = k_1 \end{cases}$$

η_k represents product k's own-price inverse elasticity of demand and γ_k represents its cross-price inverse elasticity of demand with respect to $-k$. Results will make use of the relative difference in own- and cross-price demand elasticities between products, but not absolute

try Shakeup: Q & A with Jim Handy," CBS News, 4/4/2009 http://www.cbsnews.com/news/memory-industry-shakeup-q38a-with-jim-handy/).

[12]Entry is assumed exogenous to focus on firms' strategic incentives to collude once they have entered a generation. The data show that major DRAM manufacturers generally enter each generation within one year of each other.

values of such elasticities. Firm i's profit maximization function in noncooperative play is:

$$\max_{q_{ikt}} \Pi_{it} = \sum_{t=\tau}^{\infty} \sum_{k=k_1}^{k_2} \delta^{t-\tau} \left[P\left(Q_{kt}; Q_{-kt}\right) q_{ikt} - C\left(q_{ikt}, q_{ikt-1}; c_1, \underline{c}\right) \right] \tag{1.1}$$

Firm i's first-order condition for product k_2 in period τ is:

$$\begin{aligned} \Rightarrow \frac{\partial \Pi_i}{\partial q_{ik_2\tau}} :\; P_{k_2\tau} + \frac{\partial P_{k_2\tau}}{\partial q_{ik_2\tau}} q_{ik_2\tau} &= \frac{\partial C_{ik_2\tau}}{\partial q_{ik_2\tau}} + \delta \cdot \frac{\partial C_{ik_2\tau+1}}{\partial q_{ik_2\tau}} - \frac{\partial P_{k_1\tau}}{\partial q_{ik_2\tau}} q_{ik_1\tau} \\ &= \quad c_1 \quad - \underbrace{\delta\lambda \cdot q_{ik_2\tau+1}}_{\text{Learning}} + \underbrace{\gamma_1 \cdot q_{ik_1\tau}}_{\text{Cannibalization}} \end{aligned}$$

The first-order condition illustrates the competing production incentives that firm i faces during a product's first phase. Learning-by-doing, which is based on the learning rate and discount factor, induces it to produce more than it would if cost were static.[13] It is constrained from selling output as if cost were at its end-phase minimum c_3, however, by the presence of its own competing product. The cross-price elasticity parameter γ_1 indexes the extent to which output of k_2 reduces the price of k_1.

Firms i and j play a supergame comprised of an infinite sequence of ordinary games. At the start of each period t, i's action is the value of the vector of strategic variables $s_{it} \equiv (q_{ik_1t}, q_{ik_2t})'$. The normal form strategy of the ordinary game at period t is $s_t \in \Re^2$, the vector of strategies across firms (s_{it}, s_{jt}). Per-period payoffs are $\Pi_{it}\left(s_t; \Omega\right)$, where Ω is the parameter vector. The supergame strategy of the supergame is σ_i, where:

$$\sigma_i = \begin{cases} s_{it} = f_{it}(s_1, \ldots, s_{t-1}) & \text{if } t \geq 2 \\ s_{i1} & \text{if } t = 1 \end{cases}$$

[13] To focus on the central relationship between learning and incentive compatibility, this model excludes intertemporal strategic effects of firm i's output of k in τ on firm j's output of k and $-k$ in $\{\tau + 1 \ldots \infty\}$. To the extent that firms play as intertemporal strategic substitutes, incentives to deviate are even stronger than those represented in the model. Zulehner (2003) and Siebert (2010) find evidence consistent with this phenomenon in DRAM.

Definition 1.1 *A supergame strategy is a sequence of functions which map all preceding ordinary game strategies of both players into the t^{th} ordinary game strategy of player i.*

Call the supergame strategy characterized by playing the one-shot Nash-Cournot action in each period σ_i^{NC}. A collusive equilibrium exists if firms can raise their joint profits higher than $\Pi^*\left(s_1^{NC}, s_2^{NC}, \ldots\right) \equiv \Pi^{*NC}$ by playing a supergame strategy σ_i^{-NC}. To be credible, such strategies must form a Subgame Perfect Equilibrium.[14]

Definition 1.2 *σ_i is a subgame perfect equilibrium (SPE) if the firms' strategies constitute a Nash equilibrium in every subgame of the original game.*

1.2.2 Collusion with Single-Market Punishments

To proceed, it is necessary to understand how learning and multiproduct competition shape the nature of collusion. This section begins the analysis by restricting firms to play as if k_1 and k_2 are non-interacting markets and if the only supra-competitive equilibrium that is possible is the one that maximizes the discounted sum of joint profits (i.e., the most profitable equilibrium). It compares the minimum discount factor necessary to sustain collusion for k_1 with the corresponding minimum discount factor for k_2.

Assume that the game is at period τ, where k_1 is in the second phase of its product life cycle and k_2 is in its first. Section 1.2.3 extends the analysis to the more realistic multiproduct punishment case by following the logic of Bernheim and Whinston (1990). Section 1.2.4 generalizes to allow firms to play the full range of equilibria between one-shot Nash and perfect collusion.

Consider the strategy in which each firm maximizes the discounted sum of joint profits at τ and punishes deviation on generation k at τ with Cournot reversion on k (but not $-k$) forever after.[15] Call this the joint profit maximizing collusive strategy and denote it $\sigma_i^{'}$.

[14]The equilibria discussed are not necessarily Markov Perfect only because intertemporal strategic effects are assumed away; see footnote 13. Firm i's payoff-relevant state variables at t are (1) its parametrically determined marginal cost; (2) a zero/one vector denoting firms' adherence to the collusive strategy in periods $\{1, \ldots, t-1\}$.

[15]See Abreu, 1986, 1988 for analyses of optimal trigger and stick-and-carrot punishments in supergames.

Firm i's *joint* profit maximization function at period τ is:

$$\max_{Q_{kt}} \Pi_{it} = \sum_{t=\tau}^{\infty} \sum_{k=k_1}^{k_2} \delta^{t-\tau} \left[P\left(Q_{kt}; Q_{-kt}\right) Q_{kt} - C\left(Q_{kt}, q_{ikt-1}; c_1, \underline{c}\right)\right] \tag{1.2}$$

Note that i's cost reduction for generation k is still based on its own previous-period output q_{ikt-1}, not total previous-period output Q_{kt-1}. Firms can collude in the product market but they cannot collude in the learning process to share gains from output across the two firms.[16] In contrast to noncooperative play (1.1), in cooperative play (1.2) i includes j's output decision in its own optimization problem. It therefore accounts for the strategic substitutability inherent in Cournot competition and restricts output relative to the noncooperative case.

The following definition establishes a necessary condition for σ_i' constituting a subgame perfect equilibrium.

Definition 1.3 *Define ^{*D} and ^{*J} as deviation and joint profit maximization actions of the normal form game, respectively. σ_i' is a subgame perfect equilibrium for product k at period τ if:*

$$\sum_{t=\tau+1}^{\infty} \delta^{t-\tau} \left(\Pi_{ikt}^{*J} - \Pi_{ikt}^{*D} \right) \geq \left(\Pi_{ik\tau}^{*D} - \Pi_{ik\tau}^{*J} \right)$$

The collusive strategy σ_i' must satisfy the incentive compatibility constraint at period τ. This implies that the profits from remaining on the collusive path at τ are greater than or equal to the profits from deviating at τ.

Next, I illustrate the fundamental difference between collusion on k_1 and collusion on k_2 by highlighting the difference in punishments for deviation between the two generations.

[16] Hatch and Mowery (1998) and Macher and Mowery (2003) model the different engineering processes underlying cost reduction in semiconductors and find that learning is shaped more by organizational management, data analysis and equipment technology than explicit knowledge held by individual engineers. Firms are limited in the learning they can successfully translate between plants; it is unlikely that they could increase diffusion between firms without costly effort, planning, team-oriented communication and problem solving. There is no evidence to suggest any attempt was made in the DRAM cartel. See section 1.3 for more details on the learning process.

If i deviates from k_2's collusive path in period τ fixing j's output, the profits i earns on k_2 in $\tau+1$:

$$q^{*D}_{ik_2\tau+1} = \frac{1}{3\eta_1}\left[(a - \underbrace{c_1 + \lambda q^{*D}_{ik_2\tau}}_{c_{ik_2\tau+1}}) + \underbrace{(\lambda q^{*D}_{ik_2\tau} - \lambda q^{*J}_{jk_2\tau})}_{\lambda(q^{*D}_{ik_2\tau} - q^{*J}_{jk_2\tau})} - (\gamma_1 + \gamma_2) Q_{k_1\tau+1}\right] \tag{1.3}$$

Equation 1.3 illustrates that i obtains greater market share than j on generation k_2 in period $\tau+1$ if it deviates in period τ. Deviation increases its next-period profits relative to playing a static symmetric Cournot game in a direct way and a strategic way. Directly, $c_{ik_2\tau+1} < c_{jk_2\tau+1}$ because $\lambda q^{*D}_{ik_2\tau} > \lambda q^{*J}_{jk_2\tau}$: i's second period costs are lower following deviation than following adherence because i has learned more in the first period. Consequently, $(\lambda q^{*D}_{ik_2\tau} - \lambda q^{*J}_{jk_2\tau}) < 0$: j knows that i's costs are lower, prompting j to back off and i to push forward. Firm i gains a strategic advantage over j that further increases its profits.

If i deviates on k_2 in period τ, it faces profit loss relative to foresaken collusive profits for k_2 in periods $\tau+1, \ldots, \infty$. i's foresaken profits in period $\tau+1$ are mitigated by its profit gains from the strategic advantage it generates by deviating. If i deviates on k_1 in period τ, in contrast, it forsakes future collusive profits but gains no mitigation in $\tau+1$ or any future periods. This logic leads to the following result.

Proposition 1.1 *Define δ^*_{kt} as the minimum discount factor that satisfies σ_i''s ICC at t.*
If $|\eta_1| \leq |\eta_2|$ and $|\gamma_1| \leq |\gamma_2|$ $\Rightarrow$ $\delta^*_{k_1\tau} < \delta^*_{k_2\tau}$

$$\begin{cases} \delta < \delta^*_{k_1} < \delta^*_{k_2} & \Rightarrow \textit{neither } k_1 \textit{ nor } k_2 \textit{ enforcable} \\ \delta^*_{k_1} < \delta < \delta^*_{k_2} & \Rightarrow \textit{only } k_1 \textit{ enforcable} \\ \delta^*_{k_1} < \delta^*_{k_2} < \delta & \Rightarrow \textit{both } k_1 \textit{ and } k_2 \textit{ enforcable} \end{cases}$$

Proof 1.1

See Appendix A.1.

Result 1.1 conveys the intuition that it is strictly more difficult to collude in the early phase of a product life cycle relative to a later phase. The sufficient condition establishes that as long as the same-period gain to deviating on k_1 is not substantially greater than the same-period gain to deviating on period k_2, the ICC requires a higher minimum discount factor for k_1 ($\delta^*_{k_1\tau}$) than k_2 ($\delta^*_{k_2\tau}$).[17]

The next portion of Result 1.1 depicts the relationship between the firm-invariant discount factor and the minimum discount factor necessary to enforce σ'_i. Observed conduct at τ depends on δ: if it is sufficiently or high or sufficiently low, firms are equal in their conduct toward both generations. If δ is in between the range of values, however, collusive strategy σ'_i is enforcable on k_1 but not k_2.

This result is also robust to the inclusion of learning spillovers betwen firms, as long as learning does not reduce rivals' costs more than a firm's own costs (see Appendix A.2). Section 1.2.3 considers the multiproduct punishment case, and section 1.2.4 generalizes to the full range of collusive strategies.

1.2.3 Collusion with Multiproduct Punishments

In this section I allow firms to coordinate collusive strategies across separate products linked by demand substitution.[18] Specifically, consider the strategy in which each firm maximizes the discounted sum of joint profits at τ but punishes deviation in generation k with Cournot reversion in k *and* $-k$. Call this strategy $\sigma_i^{''}$.

How will the minimum discount factor necessary to sustain collusion on generations k_1 and k_2 change when firms play $\sigma_i^{''}$ rather than $\sigma_i^{'}$? Computing equilibrium output expres-

[17]The same-period gain to deviating could be higher for k_1 than k_2 if, for example, demand is substantially less elastic for k_2 than k_1. As the opposite result is generally the case in high-technology products—where new generations spur obsolescence—this possibility is assumed away.

[18]The analysis mirrors the treatment of multimarket collusion in Bernheim and Whinston (1990) and subsequent papers.

sions shows that the disparity in incentives to deviate between products k_1 and k_2 remains. Specifically, equation 1.3 still holds: firm i's punishment on k_2 from deviating on k_2 is weakened by its learning advantage over firm j.

The difference is that i's deviation triggers punishment in period $\tau + 1$ on k_1 as well as k_2. As the marginal cost of k_1 has reached its minimum $\underline{c}$ by $\tau + 1$, foresaking collusive profits on this generation raises the punishment to deviating relative to $\sigma_i^{'}$. The following result depicts when and how $\sigma_i^{'}$ can rectify the disparity between collusion.

Proposition 1.2 *Define δ_{kt}^{**} as the minimum discount factor that satisfies $\sigma_i^{''}$'s ICC at t.*

If $|\eta_1| \leq |\eta_2|$ *and* $|\gamma_1| \leq |\gamma_2|$, $\;\Rightarrow\;\; \delta_{k_{1\tau}}^{*} = \delta_{k_{1\tau}}^{**} \;\leq \delta_{k_{2\tau}}^{**}$

Proof 1.2

See Appendix A.1.1.

Result 1.1 showed that the sustainability of collusion in single-product punishments depends on how high the discount factor is, and that there is a "window" in which the discount factor is high enough to sustain collusion on k_1 but not k_2. Result 1.2 shows that under the same conditions, multiproduct punishments are identical except for one change. The new critical discount factor for k_2 is bounded by the older critical discount factor for k_2 and the critical discount factor for k_1. Collusion is therefore at least as sustainable for k_1 as it is for k_2.

The key insight is that firms can use slack enforcement power from the incentive compatibility constraint on k_1 to enforce collusion on k_2–but only if k_1's incentive compatibility constraint has enough slack. The amount of slack depends on where the discount factor falls, so the imbalance between minimum discount thresholds necessary to sustain collusion may be unchanged, reduced, or completely eliminated. Under the same conditions as Result 1.1, collusion will always be at least as easy to sustain for k_1 as it is for k_2.

Two additional points are worth noting. First, if firms can choose between only σ_i^N and $\sigma_i^{''}$, deviation on k_2 would always imply deviation on k_1 as well. Result 1.2 highlights the fact that multiproduct punishments do not necessarily make collusion equally feasible across markets, but it is necessary to allow firms to reach a range of equilibria between the one-shot Nash and the perfectly joint profit maximizing strategies to generate more complete results. Second, although the strategies σ_i' and σ_i'' are only two among a set of infinite strategies that may support collusion at levels above the one-shot Nash equilibrium, they have the feature that punishments depend on marginal costs, which in turn depend on firms' accrued learning. Other types of trigger and stick-and-carrot punishments also possess this feature and therefore render collusion at least as sustainable for k_1 as k_2. Both of these points suggest that a fuller model of conduct will provide a testable prediction to take to the data.

1.2.4 Collusion and the Conduct Parameter

The preceding section has shown the conditions under which the minimum discount factor necessary to sustain collusion at the optimal collusive equlibrium is at least as high for the newer product at the first phase of its life cycle as the older product at its second phase. What does this result imply for observed prices and markups for firms in collusion when supra-competitive equilibria of various degrees can be reached? To answer this question, I cast the model in conjectural variation terms as follows.

Reconsider firm i's profit function and first order condition at period τ:

$$\underbrace{P_{k\tau} + \frac{\partial P_{k\tau}}{\partial q_{ik\tau}} q_{ik\tau} + \frac{\partial P_{-k\tau}}{\partial q_{ik\tau}} q_{i-k\tau}} = \frac{\partial C_{ik\tau}}{\partial q_{ik\tau}} + \delta \cdot \frac{\partial C_{ik\tau+1}}{\partial q_{ik\tau}} \tag{1.4}$$

$$\underbrace{P_{k\tau} + \underbrace{\theta_{k\tau}} \left[Q_{k\tau} \left(\frac{\partial P_{k\tau}}{\partial Q_{k\tau}} + \frac{\partial P_{-k\tau}}{\partial Q_{k\tau}} \frac{q_{i-k\tau}}{q_{ik\tau}} \right) \right]} = \qquad '' \tag{1.5}$$

$$\underbrace{\theta_{k\tau}} = \left(1 + \frac{\partial q_{jk\tau}}{\partial q_{ik\tau}} \right) \frac{q_{ik\tau}}{Q_{k\tau}} \tag{1.6}$$

In this framework, $\theta_{k\tau}$ is the conduct parameter that aggregates the average strength of collusion between firms i and j in period τ for product k. As 1.6 shows, $\theta_{k\tau}$ is based on the conjecture that i holds about j's response to a change in i's output in period t. Marginal cost is composed of the static effect and the dynamic (learning) effect as before.

Bresnahan (1989) demonstrates that the conduct parameter has a clear interpretation for several specific supply types. In Cournot equilibrium, firms in competition hold their rival's output fixed, so $\frac{\partial q_{jt}}{\partial q_{it}} = 0 \Rightarrow \theta_{kt} = \frac{1}{2}$. In the joint profit maximizing equilibrium, where firms collude perfectly, an output increase by i is met with an equal output increase by j. In that scenario, $\frac{\partial q_{jt}}{\partial q_{it}} = \frac{q_{jt}}{q_{it}} \Rightarrow \theta_{kt} = 1$. In other supra-competitive equilibria, it takes values between $\frac{1}{2}$ and 1 corresponding to the "average strength of collusion." $\theta_{kt} = \frac{3}{4}$, for example, implies that an increase in i's output is met by a corresponding (but not equal) increase in j's output.[19]

The preceding analysis suggests that firms are less likely to reach the minimum discount factor necessary to sustain collusion at τ for k_2 relative to k_1, but cannot rule out the possibility that collusion is equally successful or unsuccessful. The following result establishes a testable prediction of the hypothesis by aggregating the effectiveness of collusion into the conduct parameter.

Proposition 1.3 *If* $|\gamma_1| \leq |\gamma_2|$ *and* $|\bar{\gamma}| \leq |\gamma_2|$, $\theta_{k_1\tau} > \theta_{k_2\tau} \quad \Rightarrow \quad Q^{*J}_{k_2\tau} > Q^{*NC}_{k_2\tau}$

Proof 1.3 *See Appendix A.1.2.*

Result 1.3 provides conditions under which the possibility of collusion in both products leads demand to shift toward the newer product, k_2. If the strength of collusion, indexed by the conduct parameter, is at least as high in k_1 as k_2, and the relative demand elasticities take the same inequalities as before, then the demand curve for k_2 shifts outward because of its substitutability. Firms therefore face two competing forces in their optimization decision for k_2. Collusion induces them to restrict output via strategic substitutability, but consumer

[19] Corts (1999) shows that the empirical estimation of θ depends upon untestable functional form assumptions on the curvature of demand. I consider only θ's value in theory and do not empirically test for it.

preferences induce them to expand output. In equilibrium it is ambiguous whether i and j will jointly produce more in the collusive arrangement than they would in competition.

The result provides one further sufficient condition under which this ambiguity is resolved. There exists a cross price elasticity of demand $|\bar{\gamma}|$ such that the output expansion effect dominates the output restriction effect. This constitutes a testable prediction of the model presented thus far. Figure 1.2 depicts the result graphically.

The possibility of i and j increasing output in response to collusion changes its welfare implications in a dynamic setting, which is not captured in this model. Costs in $\tau + 1$ are $c_{i2}^{*J} < c_{i2}^{*C}$. Further conditions about the nature of competition in $\tau + 1$ in a dynamic model would establish whether or not firms pass this cost saving on to consumers.

1.3 DRAM Industry and Market

Semiconductors are crystals—usually silicon—that serve the essential role of connecting the electronic circuits that make up a microchip. Microchips are a bedrock of the electronics revolution and provide the processing, memory, speed, and performance ability of computers and many electronic devices. A critical feature of the microchip is its capacity: the number of transistors per square inch. Increases in capacity reduce the effective cost of a microchip, or, equivalently, increase its speed. Capacity has risen sharply since the industry's inception in the 1970s, famously corresponding to "Moore's Law": the maximum capacity of an individual chip doubles every 18-24 months.

Memory chips store and release data that is used in microprocessor chips such as the computer's central processing unit. DRAM operates at a specific level of the microprocessor by refreshing the transistor repeatedly. The two chips are complementary: higher processing speed is forfeited unless the device has enough memory to continuously access data and perform operations. To maximize efficiency, DRAM is sold as a package of chips attached to a circuit board to create a module. The total capacity of the chips in a mod-

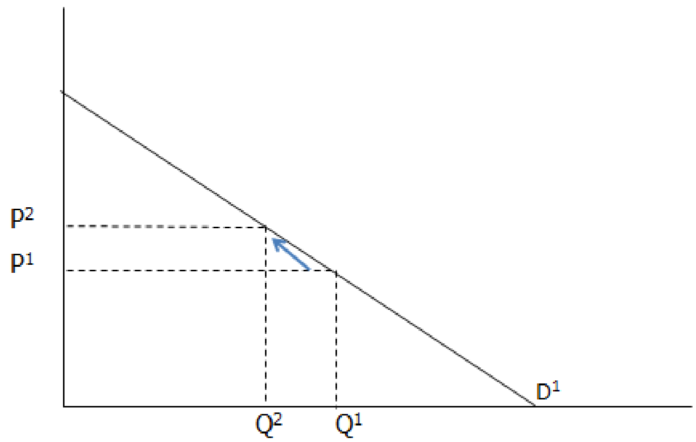

(a) Output Restriction on k_1

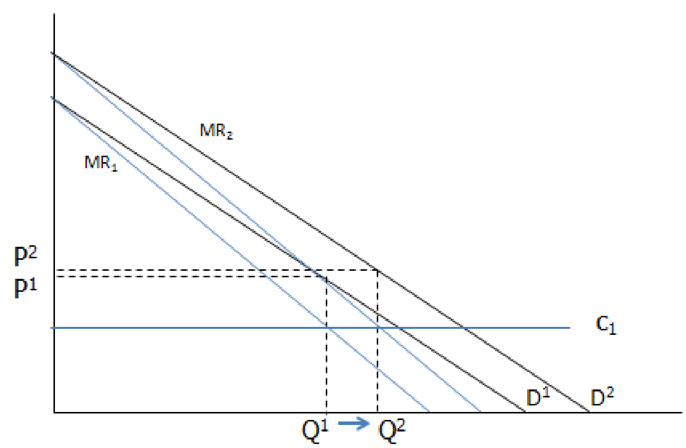

(b) Demand Shift for k_2

Figure 1.2: Collusion on k_1 Affects Demand for k_2

ule represents its *bit density* ("density"), and DRAM product generations are measured by density. Like most microchips, it is within generations that intensive learning-by-doing takes place.[20] DRAM products are largely homogeneous goods within a density generation, and are substitutable between density generations based on the customer's preferences for memory speed.[21] Several different density generations are available at any time, and because of the sequential nature of releases, they are always sequentially ordered between different points in their life cycles.

The only way to increase capacity is through *photolithography*, a fabrication process in which hundreds of chemical reactions are layered onto the underlying silicon wafer that produces microchips.[22] While finer-grain etching tools or more complex chemical reactions can advance the lithography process to increase capacity, they also require very precise light, air and dust conditions. For each duplication in capacity, the new lithographic process initially produces a yield of zero chips because such conditions are unknown. Producing a batch of chips requires hundreds of precise steps conducted at the nanoscale level. Engineers continuously fine tune chemical conditions, employ alternative techniques, phase in higher-technology equipment, and analyze the resulting data in a trial-by-error process. As Hatch and Mowery (1998) find, the most important adjustments are in *parametric processing*, determining the minute range of chemicals to be applied at each step, followed by adjustments in air particle contamination, which are impacted by the type of lithography being conducted. Problem diagnoses in these two areas gradually increase numbers of usable chips as engineers learn which factors are most conducive to functionality and obtain the technical tools to implement them. Increasing yield drives the cost of each chip down as a direct function of output, consistent with classic models of learning-by-doing.

In addition to the capacity level, DRAM innovation also occurs along the *technology*

[20] See Gruber (1998) for evidence of intergenerational learning in EPROM memory chips.

[21] Noll (2014)

[22] The other way to reduce cost per chip—but not capacity per chip—is to increase the size of the wafer itself, which allows more chips per lithographic batch. Such changes occurred about half as frequently as capacity advancements during the 1980s and 1990s; see Kang (2010) Table 3-2, citing Brown & Linden (2009).

level. Unlike the capacity innovation process, technology standards are jointly developed by fabricators with input from the largest microprocessor provider, Intel.[23] The JEDEC Solid State Technology Association ("JEDEC") is the industry's Standard Setting Organization, and it typically negotiates technology design, patents and licensing conditions several years in advance of a new rollout. In the early 2000s, two different types of new technologies were competing to become the next industry standard. In addition to JEDEC-sponsored Double Data Rate Synchronous DRAM ("DDR"), the Silicon Valley design firm Rambus developed and patented Rambus DRAM ("RDRAM"). Appendix A.3 shows that neither technology began to proliferate until after the cartel period, which rules out the possibility of competition over technologies confounding any of the empirical cartel results to follow.

The market for DRAM has greatly expanded since its start in the 1970s, primarily through increasing demand, processing speed and proliferation of computer products. DRAM chips are used directly by original equipment manufacturers (OEM's) for the memory required in personal computer (PC) desktops and notebooks, and are also purchased as standalone products to enhance the memory of existing PC's.[24] For the data period under study, PC's were the dominant application of DRAM. They were the only significant end use throughout the 1980s and early 1990s (Flamm and Reiss, 1993); they comprised about 90% of the market in the late 1990s and early 2000s (Third Amended Class Action Complaint, MDL No. 1486, pg 76); and 80% of the market in 2006 (Kang, 2010).[25]

Because of the large capital outlays for new fabrication plants and variability in demand growth for computers, the DRAM industry has been highly cyclical since its inception.[26]

[23] Source: private correspondence with Jim Handy, Director at Objective Analysis (May 19, 2014).

[24] Memory manufacturers contract with OEM's on a biweekly or monthly basis.

[25] In recent years, the market for DRAM has begun to shift from traditional PC's toward mobile smart phones and computer tablets, which use lower-power memory but sell many times the units as PC's. See "Why Growth in Mobile Devices Will Fuel Micron's DRAM Shipments," Forbes, 1/15/2013 http://www.forbes.com/sites/greatspeculations/2013/01/15/why-growth-in-mobile-devices-will-fuel-microns-dram-shipments/

[26] Capital depreciation accounted for about 50% of marginal cost in the 1990s and 80% in the 2000s. Source: see Footnote 11.

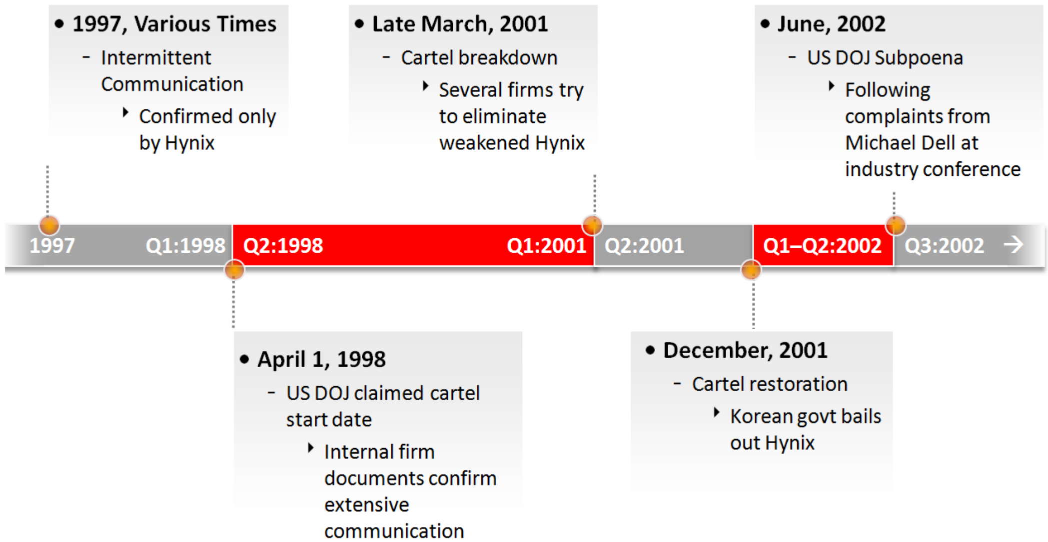

Figure 1.3: Events During Cartel Period
Sources: Noll (2014), DOJ and EC Press Releases, and US Indirect Purchaser litigation

In the late 1990s, the industry faced significant overcapacity despite steady PC growth.[27] After several firms exited the market amid falling prices and steep losses, in 1998 the largest firms formed a cartel to cut production, raise prices and restore profitability.[28]

Figure 1.3 details the events surrounding the cartel period. After intermittent communication in 1997, the cartel operated from the second quarter of 1998 through the first quarter of 2001. Noll (2014) relates that it disbanded when South Korean firm Hynix applied for bankruptcy protection, and restarted after state-controlled banks bailed out the company for $6-7 billion.[29] The cartel ended with DOJ subpoenas in June 2002 following a public accusation of collusive behavior by the CEO of Dell Inc., Michael Dell.

[27] "But too many new plants were built several years ago in a rush to benefit from a boom many thought would never end...The glut led to a free fall in prices. In-Stat expects a 21 percent decline in revenue this year, to $15.6 billion."—*Dallas Morning News*, June 1998 http://articles.chicagotribune.com/1998-06-29/business/9807100053_1_memory-in-personal-computers-dram-memory-chip-business

[28] See http://europa.eu/rapid/press-release_IP-10-586_en.htm for further information.

[29] "DRAM Rivals Question Korean Government's Role in Hynix Bailout," EE Times, 11/5/2001 http://www.eetimes.com/document.asp?doc_id=1131477.

1.4 Data

I use a proprietary dataset from industry-leading market research firm Gartner Research that lists quarterly shipments by firm and density generation, and quarterly market price by density generation, for all firms in the DRAM market from 1974-2011. I narrow the data to the period for which downstream demand data is reliably available: 1988-2011. DRAM demand is proxied by a time series of worldwide quarterly PC shipments obtained from Gartner research reports. This data is available annually from 1988-1996 and quarterly from 1997-2011.

All prices are given in US dollars and subsequently deflated to year 2000 values by the Consumer Price Index (CPI). Additionally, the Gartner data includes output per firm by DRAM technology after 2001, when it began tracking individual technology shipments within generations. Firm names allow accurate tracking of mergers, entry, and exit throughout the data sample.

In the empirical work to follow, I make use of two subtly different measures of the market price. The first (employed in previous literature) is the price *per module* of DRAM generation k sold at time t. Because DRAM chips are sold as modules, it is the actual price that buyers pay for a given DRAM generation. Another way to represent DRAM price is through the price of an individual chip standardized by its capacity. The DRAM price *per MB* is the market-wide price of a chip divided by its number of megabytes (MB's). Both price definitions are equivalent.

Figure 1.4 plots the price per module by generation for all 11 generations of DRAM from 1988-2011. The pattern is clear: price is high in initial periods, and it declines in a logarithmic shape as a generation ages. The most significant price changes occur within the first five to ten years of a generation. This is consistent with classic models of learning-by-doing.

Figure 1.5 shows the price per MB averaged and weighted by sales across generations, from 1988-2011. The standardized price declined through the sample because a

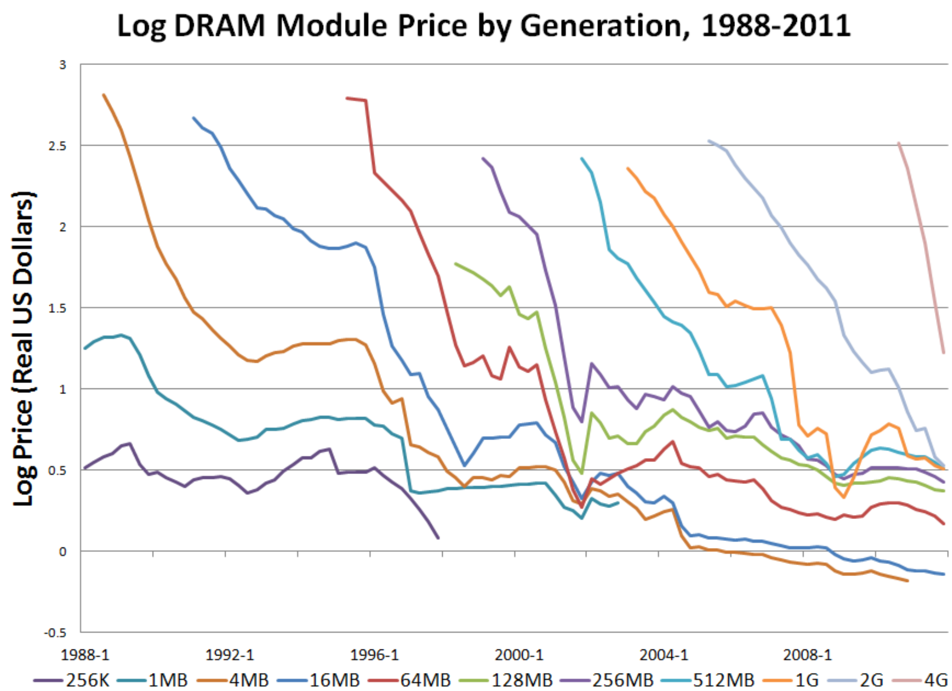

Figure 1.4: DRAM Module Log Price Trends By Generation

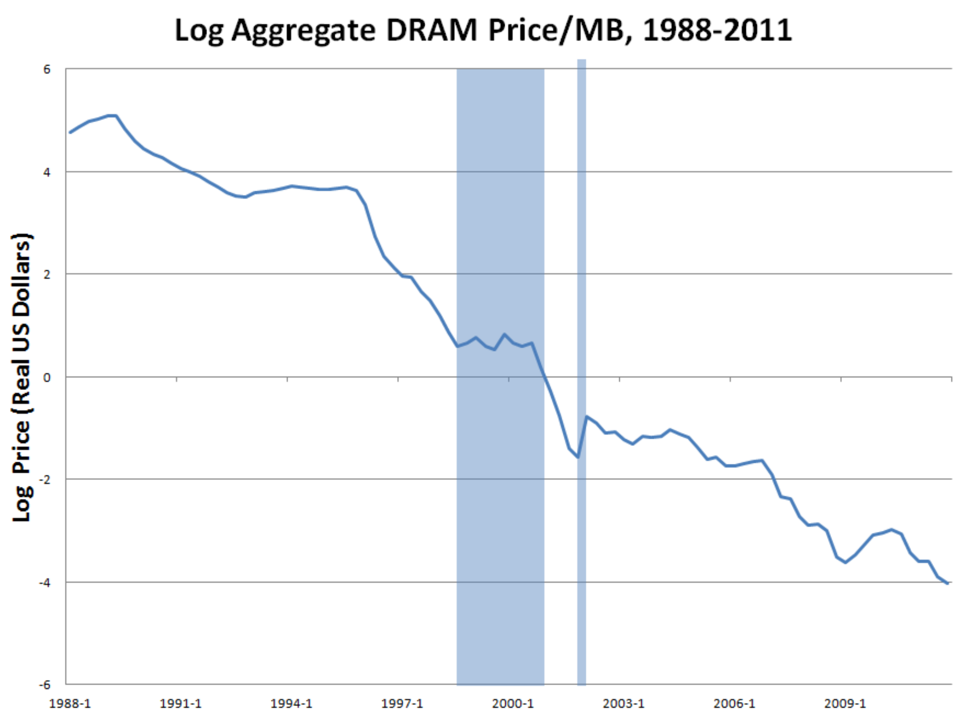

Figure 1.5: Overall DRAM Price Trend with Cartel Periods

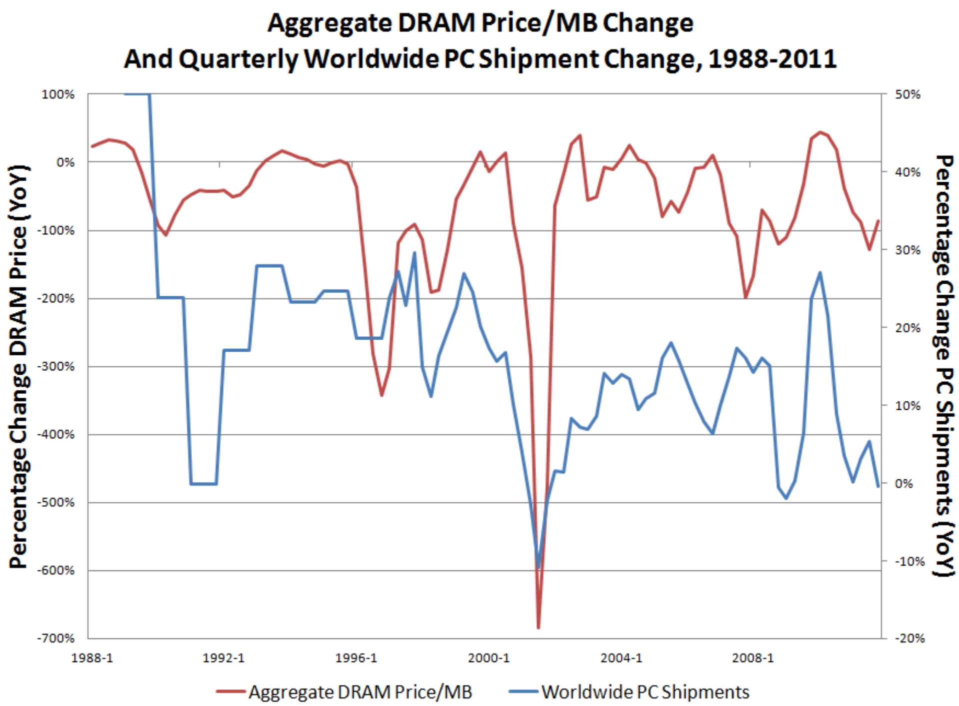

Figure 1.6: Aggregate DRAM Price Change against PC Shipment Change

fixed DRAM chip becomes cheaper as density increases. However, there were multiyear cycles of sharp price decreases interspersed with price plateaus or slight increases. The shaded area represents periods of cartel activity. The figure shows that price vacillated during the cartel period, staying roughly constant from 1998-2000 before dropping in 2001. The drop coincides with the period at which the dot-com bubble burst and the cartel disbanded. Price increased in the first two quarters of 2002, when the cartel regrouped, before declining again after mid-2002 as the DOJ initiated its investigation.

Figure 1.6 plots the change in aggregate DRAM price per MB against the change in worldwide PC shipments. It shows that most of the variation in DRAM price can be explained by variation in PC shipments. The effects of the 2001 dot-com bubble crash are particularly striking: both series reach their troughs in concert. Regression specifications make use of the change in PC shipment rates in explaining generation-level DRAM price throughout the product life cycle.

Forty-eight different firms appear in the dataset, with a maximum of 24 active firms in

1996.[30] As many as seven different product generations were active at any given cross-sectional point in the dataset, although two to three accounted for most of the sales.[31] Eleven different product generations, from 256Kb to 2Gb, appear with sufficient frequency to include in the analysis.[32] Table A.1 lists each DRAM generation and its years active within the sample period. At the market level, the industry is a classic oligopoly: ten firms accounted for 85% of output across generations in 2000. The four largest firms—Samsung, Infineon, Micron, and Hynix—held 63% of the market.[33] See Table A.2 for C_4 concentration ratios by generation from 1994-2004.

1.5 Empirical Results

The theoretical model presented above generates testable predictions for the effect of collusion on market prices and outcomes. It is based on the intuition that collusive equilibria are at least as difficult to sustain on new generations as old generations. The larger is the disparity in average strength of collusion between generations, the larger is the demand shift toward newer generations. This leads to a sufficient condition that is testable: newer generations during the cartel period sell more output than they would in competition. It also implies that the cartel period overcharge is higher for older generations than newer generations, and ambiguous in sign for newer generations.

I provide two types of estimates: (1) the effect of collusion on the dependent variable, output or price; (2) the effect of collusion on the dependent variable with respect to product cycle age. The first relationship is identified by repeated product life cycles among generations before, during, and after the cartel. I compare observations a given amount

[30] See Siebert and Zulehner (2013) for an empirical analysis of firm entry and shakeout in the DRAM industry.

[31] See fig. A.3 for the evolution in market shares between five DRAM generations active during the cartel period.

[32] I also discard two minor "mini-generations," 2Mb and 8Mb, that briefly appear between industry standard generations in the 1990s.

[33] I pool the sales of LG Semiconductor with Hynix, which purchased the company in 1999.

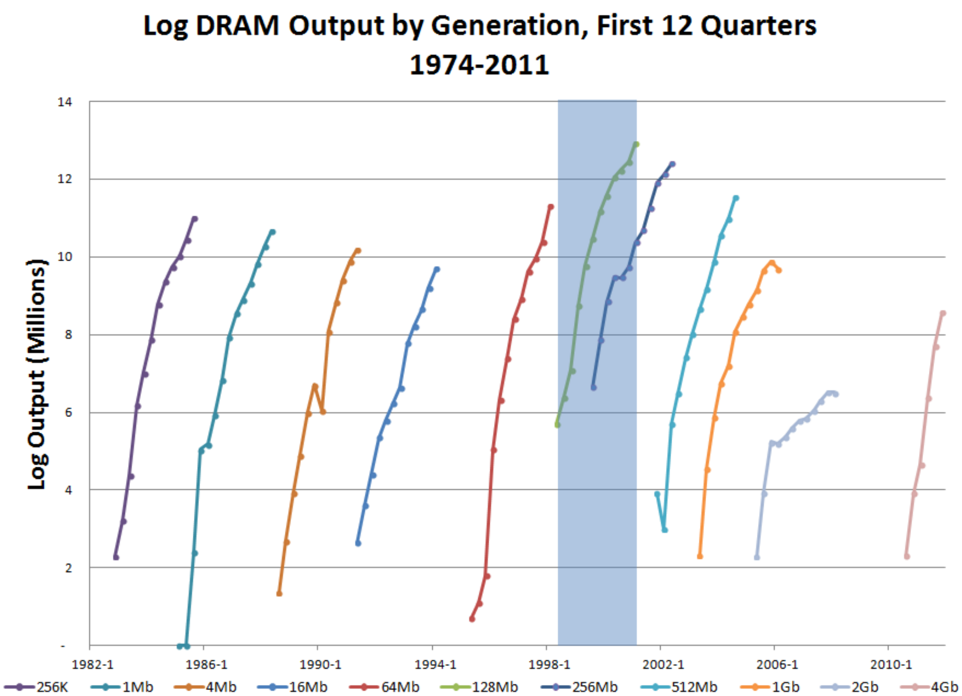

Figure 1.7: DRAM Output by Generation, First 12 Quarters

of quarters from generation or firm entry *within* the cartel period to other observations the same amount of quarters from generation or firm entry *outside* the cartel period, adding a time-varying proxy for demand.

The second relationship is the differential effect of collusion with respect to generation age. Its identification arises from two sources. First, because the cartel was active for 12 quarters, each active generation provides time-series variation in age. Second, five different generations were at sequentially ordered stages of their product cycle at each point in time during the cartel. This provides cross-sectional variation between generations. The latter source is especially strong in ruling out the possibility that changes in firms' discount factors over time—through macroeconomic conditions, firm-specific financial health or management changes—bias the results.

1.5.1 Cartel Output Effects

Figure 1.7 plots logged generation-level output of DRAM for the first 12 quarters of each generation. It displays data from the 11 even generations under study, with each dot representing one quarter of output. Output generally increases from quarter to quarter during the early stage of production, as existing firms increase their output and new firms enter.

Firms started selling units of 128Mb and 256Mb chips during the cartel period, which is shaded. 128Mb was the newest generation from 1998 until 1999, when firms began producing 256Mb chips. These two generations were therefore in the early stage of their product cycles during the cartel period. Sales of the 128Mb and 256Mb generations immediately stand out: initial output values are orders of magnitude greater than the corresponding values during non-cartel periods before and after the conspiracy.[34] It is especially noteworthy that 128Mb and 256Mb output is greater than post-cartel generations, because the increasing market for DRAM over time implies gradually increasing overall output.[35]

To further investigate the hypothesis that cartel output increased for the newer generations, I estimate the following regression:

$$\begin{aligned} Log(q_{ikt}) = \beta_1 FirmAge_{ikt} + \beta_2 FirmAge_{ikt}^2 + \beta_3 GenAge_{kt} + \beta_4 GenAge_{kt}^2 \\ + \beta_5 Y_t + \beta_6 Y_t \times GenAge_{kt} + \beta_7 Y_t \times GenAge_{kt}^2 + \beta_k \mathbb{1}(Coll)_{kt} + \epsilon_{ikt} \end{aligned} \tag{1.7}$$

This specification regresses the log of a firm's output for generation k at time t on measures of firm age, generation age, demand, and an indicator variable for the cartel period. $FirmAge_{ikt}$ and its square denote the elapsed quarters since firm i began production in generation k. $GenAge_{kt}$ and its square represent the elapsed quarters since the first firm in generation k began output. These variables account for the positive trend in output during

[34] Initial output values for the three pre-1982 generations, 4Kb, 16Kb and 64Kb, are also far lower than those for 128Mb and 256Mb.

[35] Figure A.5 shows that peak-stage output indeed increases greatly between generations over time.

the early phase of a product cycle, when cost decreases and firms pursue learning gains, to the negative trend later in the product life cycle, when output crosses its peak and firms shift production to a newer generation. $\mathbb{1}(Coll)_{kt}$ is an indicator variable equal to one during generation k's cartel period and zero otherwise. This variable permits separate estimation of the effect of cartelization on output by generation.

Y_t is a demand shifter representing the growth rate of worldwide PC shipments. Departures from the baseline rate of PC growth, which is mostly positive through the sample period, are taken as exogeneous changes in the demand for DRAM.[36] Interactions between Y_t and $GenAge_{jt}$ allow the effect of a change in PC shipments to be stronger in the initial phase of a product cycle, when OEM's may increase production of memory-intensive computers, than later stages.

Table 1.1 shows that the robust industry-level effect observed in Figure 1.7 carries over to the firm level as well. It displays the results of regression 1.7 with and without controls for demand. The coefficients for firm and generation age take the expected signs, with primary terms positive and squared terms negative signaling the pre- and post-peak phases of the product life cycle. The demand proxies in the right hand side panel also take the expected signs, with the PC Growth Rate positively associated with output levels at the earliest stages of the product cycle.

The cartel binary variables are strongly positive and consistent for the 64Mb, 128Mb and 256Mb generations in both specifications, but negative in the 16Mb generation and negatively significant in the 4Mb generation. Wald tests for equality reject the joint hypotheses that either of the latter two coefficients arise from the same distribution as either of the former three coefficients. These results provide strong evidence that the cartel's effect was twofold: participants restricted output on older generations and raised output on newer generations.

[36]The proliferation of PC's from the 1980s through the 2000s was due to the decreasing overall cost of the computer's components, of which DRAM only comprises 5-10%. There is no evidence that DRAM prices themselves drive a significant portion of PC sales. See fig. 1.6 for the correlation between Y_t and the aggregate price of DRAM.

Table 1.1: Log Firm Output on Covariates, 1988-2011

	(1)		(2)	
	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
Firm Age	0.138***	(0.007)	0.143***	(0.007)
Firm Age^2	-0.002***	(0.001)	-0.002***	(0.001)
Gen Age	0.347***	(0.005)	0.348***	(0.006)
Gen Age^2	-0.004***	(0.001)	-0.004***	(0.001)
PC Growth Rate			0.256***	(0.008)
PC Growth Rate $\times$ Gen Age			-0.013***	(0.001)
PC Growth Rate $\times$ Gen Age^2			0.001**	(0.000)
I(Coll)$_4$	-0.379***	(0.101)	-0.314***	(0.089)
I(Coll)$_{16}$	-0.117	(0.079)	-0.058	(0.073)
I(Coll)$_{64}$	3.122***	(0.191)	2.141***	(0.174)
I(Coll)$_{128}$	4.852***	(0.257)	2.569***	(0.241)
I(Coll)$_{256}$	3.743***	(0.365)	1.628***	(0.335)
R^2	0.914		0.9273	
N	6640		6640	

***$p < 0.01$, **$p < 0.05$, *$p < 0.1$. The dependent variable is logged firm-level output of DRAM units, by quarter. Firm- and generation-age represent the elapsed quarters since the first firm in generation k began output. I(Coll)$_k$ estimates the average change in output for generation k during the cartel period.

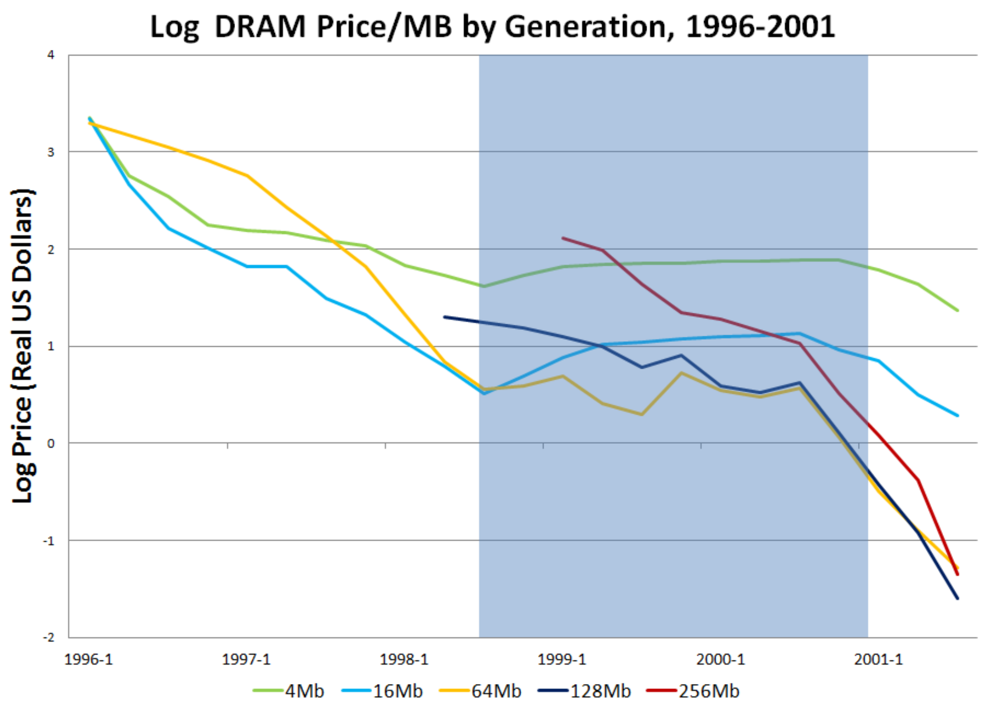

Figure 1.8: DRAM Price/MB Trend During First Cartel Period

1.5.2 Cartel Price Effects

I now describe and estimate the effect of the DRAM cartel on the price of each active generation. Figure 1.8 uses DRAM prices per MB to represent all five cartel generations at the same time on the same axis. The shaded area depicts the cartel period. By inspection, the price of the older, outgoing 4Mb and 16Mb chips steadily rise, while 64Mb price rises and falls jaggedly to stay roughly equal on average. The prices of the newer two chips continue to fall.

Figure 1.9 restricts the price series to the 128Mb and 256Mb generations, which entered during the first period of collusion and accounted for the majority of market share by the second period. It extends the x-axis to show the price trend within these generations over the first and second periods of collusion. Prices decline in the first phase of collusion but rise sharply in the second phase before the cartel is disbanded. The visual evidence in figures 1.8 and 1.9 is consistent with the results of the model: the effect of collusion on price appears to be larger as the product cycle progresses.

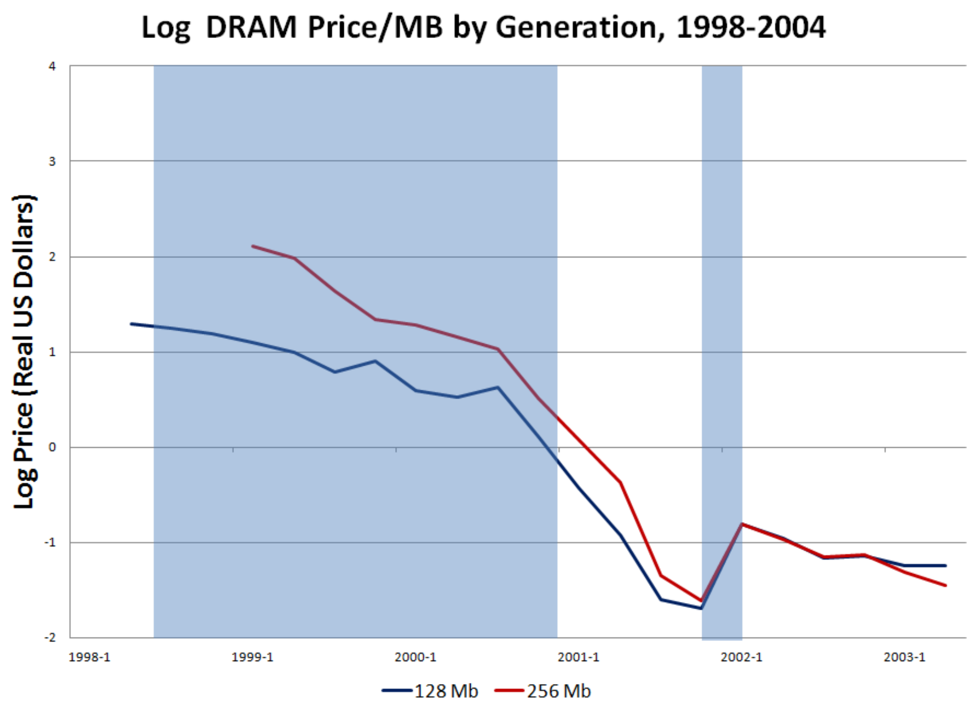

Figure 1.9: DRAM Price/MB Trend From First to Second Cartel Period

I test this intuition more formally by controlling for the product cycle phase in the following regression.

$$
\begin{aligned}
ln(P_{jt}) = \alpha + \beta_1 LogGenAge_{jt} + \beta_2\, Y_t + \beta_3\, Y_t \times GenAge_{jt} + \beta_4\, Y_t \times GenAge^2_{jt} \\
+ \beta_5 \mathbb{1}(Coll) + \epsilon_{jt}
\end{aligned}
\tag{1.8}
$$

This regression differs from 1.7 in two ways. First, the dependent variable $ln(P_{jt})$ is the logged industry-level price of generation k at time t. Second, the regression is run separately for each of the five cartel generations. The sample for each regression is restricted to a total of 24 quarters before, during, and after the quarters of cartel activity for the cartel generation in question. This procedure is conducted to maximize the statistical power of the test because all variation is at the generation-quarter level. $\mathbb{1}(Coll)$ therefore estimates the average effect of the cartel on the price of k.

Table 1.2: Log Module Price Level on Covariates, 1988-2011

	4Mb		16Mb	
Intercept	6.360***	5.088***	5.316***	2.841***
	(0.501)	(0.833)	(0.496)	(0.757)
Log Gen Age	-1.551***	-1.218***	-1.301***	-0.597***
	(0.131)	(0.218)	(0.141)	(0.215)
PC Growth Rate	0.502**	10.214	1.488***	13.482**
	(0.234)	(7.124)	(0.141)	(5.917)
PC Growth Rate $\times$ Gen Age		-0.367		-0.527
		(0.313)		(0.336)
PC Growth Rate $\times$ Gen Age^2		0.003		-0.005
		(0.003)		(.005)
I(Collusion)	0.067	0.081	0.243**	0.262**
	(0.077)	(0.078)	(0.122)	(0.120)
Quarters	34-57	34-57	24-47	24-47
N	173	173	200	200
R^2	0.477	0.490	0.350	0.405

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The dependent variable is logged industry-level deflated price in USD, by generation and quarter. The indicator variable I(Collusion) estimates the average price change in the given generation during the cartel period.

Table 1.3: Log Module Price Level on Covariates, 1988-2011

	64Mb		128Mb		256Mb	
Intercept	7.642***	7.033***	5.651***	5.524***	5.629***	5.494***
	(0.298)	(0.553)	(0.234)	(0.172)	(0.173)	(0.233)
Log Gen Age	-2.091***	-1.882***	-1.353***	-1.280***	-1.335***	-1.258***
	(0.100)	(0.189)	(0.066)	(0.099)	(0.066)	(0.100)
PC Growth Rate	3.903***	4.933*	4.160***	1.865	4.030***	1.740
	(0.477)	(2.965)	(0.617)	(1.676)	(0.621)	(1.660)
PC Growth Rate $\times$ Gen Age		0.024		0.710***		0.732***
		(0.291)		(0.269)		(0.267)
PC Growth Rate $\times$ Gen Age^2		-0.003		-.034***		-.036***
		(0.007)		(0.010)		(0.010)
I(Collusion)	-0.263*	-0.289*	-0.644***	-0.701***	-0.294	-0.260
	(0.160)	(0.160)	(0.197)	(0.190)	(0.208)	(0.199)
2-Sided Eq. Test w/ 4Mb		(**)		(***)		()
2-Sided Eq. Test w/ 16Mb		(***)		(***)		(***)
2-Sided Eq. Test w/ 64Mb				(*)		()
Quarters	7-30	7-30	1-23	1-23	1-23	1-23
N	201	201	178	178	178	178
R^2	0.718	0.722	0.720	0.748	0.720	0.758

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The dependent variable is logged industry-level deflated price in USD, by generation and quarter. The indicator variable I(Collusion) estimates the average price change in the given generation during the cartel period.

Table 1.2 and table 1.3 show the results from regression 1.8 with and without interactions of DRAM demand over the product cycle. Generation age enters negatively and significantly in all specifications, consistent with a decreasing effect of the product cycle phase on price. The coefficients for Y_t and its interactions indicate that the effect of PC growth on DRAM price is positive and significant, that it increases in magnitude in the ramp-up phase of a product cycle, and that its effect attenuates as a cycle passes its peak.

The second-to-bottom row in all specifications highlights the average effect of collusion on price for generation k. The estimates are consistent with the model's predictions: the average effect of collusion on price appears positively for the 4Mb and 16Mb generations, but significantly lower for the 64Mb, 128Mb and 256Mb generations.[37]

In fact, while the collusion coefficient is about 25% and significantly different from zero for the 16Mb generation, it is negative and significantly different from zero in the 64Mb, 128Mb and 256Mb generations.[38] The magnitudes of the price decreases are particularly noteworthy: the 128Mb generation, for example, is estimated to price 70% lower during the cartel than the competitive counterfactual. The bottom row of table 1.3 displays the p-value from a joint test of equality between the collusion coefficient for each generation and its follower. Equality is rejected strongly for 128Mb, while the p-value ranges from about 10-20% for 64Mb and 256Mb.

These findings lend further credence to table 1.1's result that successful collusion on older DRAM generations shifted demand to newer generations, which produced more during the DRAM cartel than they would have in competition. They imply that increased output created increased overall learning, and that firms passed on a component of cost savings to consumers in the 64Mb, 128Mb, and 256Mb generations.

Regression 1.8 identifies the effect of collusion on price for generation k under the fol-

[37]Figure A.4 shows that while the share of generation market share that cartel members comprised was relatively constant and at least 80% for the 16Mb through 256Mb generations, it was only 50% and declining for the 4Mb generation. Insider market share declined as the largest firms exited the generation and smaller fringe firms entered. It is therefore unclear whether to expect a positive or zero effect on cartel price for the 4Mb generation.

[38]The 256Mb generation entered in Q1-1999. It therefore has only nine quarters of cartel activity.

lowing assumptions: (1) PC shipment growth affects k and $-k$'s price equally at equal generation age; (2) supply-side factors vary as a function of generation age and not time, age-time or age-firm-time. Examples of violations of (2) include differential levels of capacity utilization between periods (systematic over- or under-capacity) or tacit collusion in non-cartel periods. If industry-wide changes in supply-side behavior vary over time—and not over the product cycle—they bias the $\mathbb{1}(Collusion)$ coefficient equally in all five regressions. Because multiple generations are always on the market at the same time, industry-wide changes apply to generations at the beginning, middle and end of their product cycles simultaneously.

Regressions of average price level do not capture any time trend in price during the cartel period. If the rate of price overcharge increases as a function of product age, as the model predicts it does, then the average price level understates the price effect of collusion. To account for these possibilities, I add a cartel indicator-age interaction term to regression 1.8 and display the results below.

Table 1.4 displays empirical results from the specification in which generations are pooled and the cartel indicator averages across all cartel generations. Results from specifications with and without the 128Mb generation are consistent with the model's predictions: the interacted coefficient of cartelization and generation age is strongly positive and significant. This indicates that the effect of collusion on prices increases as a function of generation age.

Figure 1.10 plots predicted price coefficients from the model with and without the binary variable indicating cartel activity. The x-axis is generation age and the y-axis is logged price per module. The figure estimates the expected price path of a generation beginning collusion from its first period through the rest of its product cycle. It is consistent with a causal link between collusion in one generation and demand shift to a preceding generation. Specifically, learning effects through output expansion push the price below its competitive counterpart for the first 20-24 quarters, while output restriction raises the price

Table 1.4: Log Module Price Trend on Covariates, 1988-2011

	(1) All Generations	(2) 128Mb Omitted
Intercept	5.639***	5.906***
	(0.117)	(0.116)
Log Gen Age	-1.382***	-1.450***
	(0.034)	(0.034)
PC Growth Rate	5.910***	5.547***
	(0.584)	(0.573)
PC Growth Rate $\times$ Gen Age	-0.141***	-0.136***
	(0.029)	(0.028)
PC Growth Rate $\times$ Gen Age^2	0.0006*	0.001**
	(0.0003)	(0.0003)
I(Coll)	-1.080***	-0.912***
	(0.200)	(0.216)
I(Coll) $\times$ Log Gen Age	0.334***	0.294***
	(0.062)	(0.065)
N	580	525
R^2	0.907	0.921

***$p < 0.01$, **$p < 0.05$, *$p < 0.1$. The dependent variable is logged industry-level deflated price in USD, by generation and quarter. PC Growth Rate is a demand proxy that measures the year-over-year change in worldwide PC shipments. Interactions of PC Growth Rate capture the differential effect of a change in shipment rates on DRAM price over the early, peak, and post-peak phases of the product cycle. The indicator variable I(Collusion) estimates the average price change from cartel activity at the first period of a generation's product cycle. I(Collusion) $\times$ Log Gen Age estimates the price trend from cartel activity during a generation's product cycle.

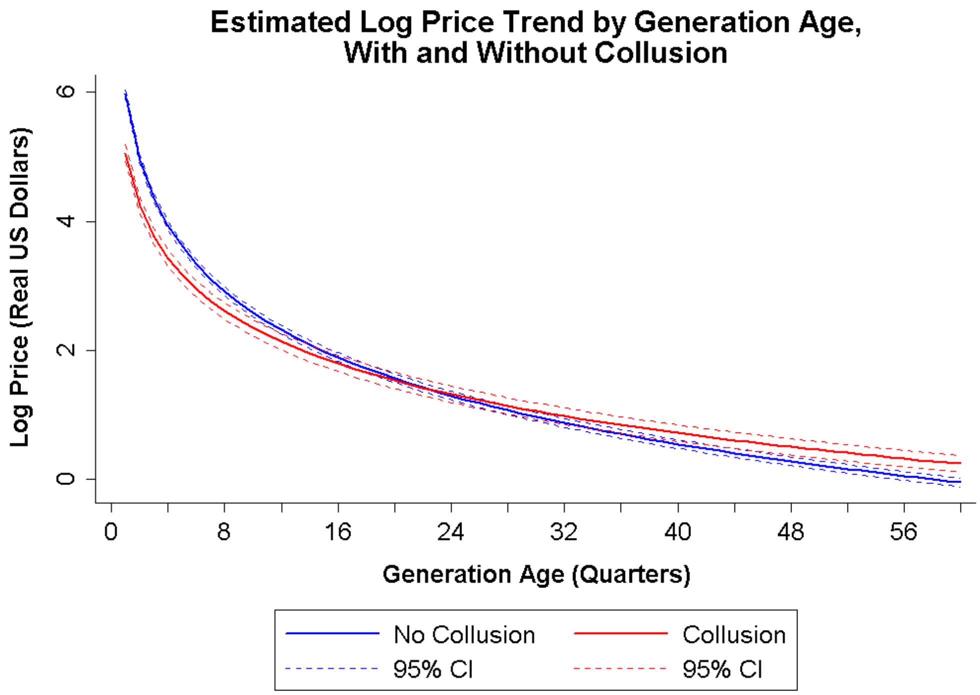

Figure 1.10: Estimated Competitive and Collusive Price Trends

in the following quarters. The break-even point occurs roughly at a generation's peak in its sixth year.

1.6 Conclusion

If firms in learning-by-doing industries refrain from competition, the cost to society could be especially large, because firms also reduce the rate at which they lower long-term costs. The fundamental insight of this paper is that the effectiveness of collusion in such markets is determined by the rate of learning. The effectiveness of collusion in one product generation impacts equilibrium price and output of competing product generations. In particular, successful collusion in a mature generation shifts demand toward newer product generations. The shift renders collusion's effect on output ambiguous for newer generations: coordination induces them to cut output, while the increase in demand induces them to raise output. Increased net output results in *more* learning in collusion than competition for newer generations.

Empirical analysis of the DRAM cartel shows evidence consistent with the model. Firms are estimated to produce substantially more output for the frontier 64Mb, 128Mb and 256Mb generations during collusion than competiton but less output for the established 4Mb and 16Mb generations. Most interestingly, prices of the three frontier generations are estimated to be significantly lower as a result of collusion. This implies that, consistent with the model, firms learned more during 12 quarters of collusion than they would have during 12 quarters of competition.

It is also noteworthy that these results are not unique to the DRAM market or its episode of collusion. In the aftermath of the DRAM cartel, many of the largest electronics manufacturers in the world have settled charges of price fixing in related markets, including LCD panels and hard disk drives. There is strong evidence that LCD makers, spurred by increased incentives to roll out newer generations, increased the pace of these introductions during collusion relative to competition. Further research and study of different markets should be conducted to reveal the extent of the substitution and its welfare effects.

This paper carries several implications for antitrust policy, particularly as it relates to many high-technology industries that experience large cost reductions through learning. Economists have long recognized that such industries may under- or over-invest in output early in the product cycle depending on the learning rate, the extent of interfirm learning spillovers, and strategic entry deterrence. The present study suggests that if firms receive a subsidy or other credit designed to raise their early stage output, they have strong incentives to pass on the resulting cost savings accrued by learning. These incentives can remain economically significant even when firms are acting as a cartel.

It further suggests to competition authorities monitoring high-profile, learning-intensive markets that collusion does not always imply increased prices on all generations of a product. Indeed, faster-than-usual price decreases may signal collusion among older generations. Moreover, as technology becomes increasingly expensive to downsize at the nanoscale level, it is well understood that semiconductors of all types are approaching a limit to cost reduction.[39] This means that overall learning rates are slowing down while other structural features that facilitate coordination—product homogeneity, multimarket contact, R&D cooperation—remain in place. Competition authorities should be aware that the conditions favoring shakeout may also raise the likelihood of firms reaching collusive equilibria, as they did in this case.

[39] Mark Bohr, senior fellow and director at Intel, relates: "Everybody in our industry will acknowledge it is getting tougher with every new generation, [but] we are going to carry the Moore's Law banner as far as we can."—"Intel Details 14-Nanometer Chip Aimed at Tablets," *Wall Street Journal* 8/11/2014. http://online.wsj.com/articles/intel-details-new-chip-aimed-at-tablets-1407775008

CHAPTER 2

Tort Liability and Settlement Failure: Evidence on Litigated Auto Insurance Claims

(With Sharon Tennyson, Department of Policy Analysis & Management, Cornell University)

This paper empirically tests the predictions of the Priest-Klein model of pre-trial bargaining. It exploits variation in tort liability for bad faith insurance law across states and time during two decades of evolving law from the 1970s to the 1990s. Using repeated cross-sectional datasets of auto insurance claims from the Insurance Research Council, it finds evidence consistent with the hypothesis that variance in parties' subjective estimates of trial outcomes drove the likelihood of settlement. The likelihood of trial for an average claim is estimated to have risen by over 20% in the initial years following reform among the first group of states to enact the tort remedy. Trial rates among tort states thereafter declined through the sample, dropping over 10% below control states by 1997. A similar relationship is estimated for the likelihood of a lawsuit being filed, and characteristics of litigated claims are consistent with a different subset of claims being disputed following regime change. Results are robust to sample selection bias, endogeneity in settlement time, and other state-level legislation on punitive damages limits and prejudgment interest. While there is limited evidence for the predictions of asymmetric information models of settlement, we conclude that policyholders and insurers negotiated in a manner consistent with divergent expectations.

2.1 Introduction

The theory of law and economics has long emphasized the link between legal rights and settlement bargaining. It has shown that the welfare effects of rights, such as those granted under tort law, depend critically on the comparative statics of pre-trial bargaining. Among the most important of these are the likelihood that a dispute reaches litigation and the likelihoood that litigation reaches trial.[1] It is well known, for example, that civil case backlog at the county, state and federal levels imposes costly strain on the American legal system.[2]

But there remain different frameworks to analyze breakdown in pre-trial bargaining. Divergent expectations models of litigation posit that parties hold mutually optimistic assessments of their chances at trial.[3] Parties fail to settle if the litigation costs of trial are small enough to justify the expected utility gain of trial relative to settlement. Asymmetric information models, in contrast, assume that one party holds an objectively correct prior assessment of the expected trial judgment.[4] The other party knows only the distribution of priors. If the plaintiff is uninformed, she may demand a settlement amount that is greater than the expected cost the defendant would expect to pay at trial, which causes disagreement.[5]

In this paper, we present an empirical analysis of settlement behavior following a change in the body of law and show that it is consistent with the divergent expectations model proposed by Priest and Klein (1984). We focus on the evolution of insurance "bad faith" law, one of the most active areas of civil liability expansion in the United States over

[1]Throughout the paper, we use "litigation" to mean the presence of a suit filed by plaintiff against defendant.

[2]See http://www.wsj.com/articles/in-federal-courts-civil-cases-pile-up-1428343746 for a recent description of the causes and consequences of federal case backlog.

[3]These models date to the early literature on law and economics; see Landes (1971), Gould (1973), Posner (1973), Landes and Posner (1979), and Shavell (1982).

[4]Seminal models based on screening between informed types include Bebchuk (1984), P'ng (1983), and Spier (1992). Seminal models of signaling between plaintiff and defendant include Reinganum and Wilde (1986) and Nalebuff (1987). See Spier (2007) for a review of asymmetric information models.

[5]The comparative statics of settlement probability in such models are typically invariant to whether the plaintiff or the defendant has the informational advantage.

the closing decades of the twentieth century. Insurance law requires providers to deal fairly, thoroughly, and promptly in settling insurance claims with policyholders. Breach of this duty constitutes a bad faith liability. Policyholders who dispute the value of a claim can sue their insurer for underpayment, and also for bad faith negotiation.

Specifically, we estimate how the tort remedy for first-party bad faith changed the incidence of settlement failure between automobile insurers and their policyholders.[6] We exploit three primary sources of explanatory power to test the model. First, states differ in whether they recognize a cause of action for the tort of bad faith. This creates two groups of states that differ in two ways: the standard necessary to establish bad faith and the magnitude of damages once it has been established. In New York, for example, common law takes the traditional position that bad faith dealings are a part of contract law.[7] Because the state does not recognize bad faith as an independent cause of action, liability would typically entitle plaintiffs to the compensatory value of the claim, but no extra-contractual or noneconomic damages. In Wisconsin, however, common law recognizes bad faith and adjudicates it as a tortious harm.[8] To establish bad faith liability in Wisconsin, plaintiffs must convince the court that the insurer's behavior meets elements of both an intentional and negligence standard.[9] A liability finding would grant plaintiffs compensatory damages, as well as consideration for additional awards: economic harm above the policy limits, emotional distress, and expressly punitive damages.[10]

Second, states changed bad faith regimes at different points in time, and state-level

[6]This paper's exclusive focus is on first-party rather than third-party bad faith insurance. Bad faith is used to mean "first-party" throughout the paper, unless otherwise noted. See section 2.3 for the distinction between the two types of insurance.

[7]*Gordon v. Nationwide Mut. Ins. Co.*, 30 NY. 2d 427 (1972).

[8]*Anderson v. Continental Insurance Co.*, 85 Wis.2d 675, 271 N.W.2d 368 (1978). Most states that recognize a private cause of action for first-party bad faith adjudicate it through tort law. Several states do recognize the cause of action but adjudicate it through contract or statute; see further details in section 2.3, and a wider discussion in Tennyson and Warfel (2010).

[9]See Sykes (1996), pp. 411-412. The *Gordon* standard is the most common model for first-party bad faith around the country.

[10]Punitive damages have been granted in a number of such cases and follow naturally from the tort of bad faith. As Sykes (1996) explains: "once 'bad faith' is established under the applicable standard, the full array of tort remedies generally becomes available...[and] because 'bad faith' can often be characterized as 'reckless,' 'malicious,' 'wanton,' and the like, many plaintiffs will be able to collect punitive damages."

reform has undergone a relatively discrete start and finish. The Priest-Klein model distinguishes between transitory and permanent levels of uncertainty in beliefs about trial success. To avoid conflating the two, it is necessary to estimate how comparative statics change over time. It is also necessary for the degree of liability and evidentiary standards within states to stabilize, so that it is possible to distinguish short- from long-term effects. The study of bad faith law makes such an analysis possible: states began determining adjudication for first-party bad faith following a landmark 1973 California decision.[11] In the subsequent two decades, the majority of states followed California by granting a tort-based cause of action for bad faith, while others clarified that it did not exist. After the early 1990s, only a handful of states would modify existing bad faith regimes significantly.[12]

Finally, auto insurance claims provide a set of frequent, high-information bargaining events against the backdrop of a trial. Changes in bad faith insurance law are unlikely to induce significant changes in the primary incentives to file a claim. Because the law pertains to an insurer's liability for bad faith, and not more generally to an individual's liability for personal injury or property, a bad faith tort can occur only *after* a claim has been filed. This stands in sharp contrast to more widely studied applications of tort law: auto accidents, medical malpractice, and products liability.[13] We present a model that builds on these three institutional features of bad faith insurance: variation in state law, variation over time, and the distribution of underlying claims between the two legal regimes. The model develops conditions under which a sample of insurance claims from states with and without tort liability can identify the effect of the law on settlement failure probabilities.

Such a sample is employed with the Insurance Research Council's national decennial survey of auto insurance claims. We build a repeated cross-sectional dataset with survey

[11] *Gruenberg v. Aetna Ins. Co.*, 510 P.2d 1032, 1037-38 (Cal. 1973). *Gruenberg*, the other primary model that states use for first-party bad faith, is commonly understood to be less narrow than *Gordon*.

[12] There are two primary reasons. Where states treated bad faith as common law, judicial precedent clarified the scope and limitations of the law and increased the predictability of judgments. In addition, nationwide doctrinal debates surrounding tort and medical malpractice liability spurred courts to examine bad faith during the 1960s and 1970s. This movement slowed down by the 1990s. See Abraham (1994) for a full analysis of the maturity of bad faith law.

[13] See Kessler and Rubinfeld (2007) for a survey of these literatures.

responses from 1977, 1987, and 1997. We exploit natural variation in state bad faith regime and time to directly identify the effect of the law on the probability of settlement failure. We examine binary measures of settlement failure at two different stages of the negotiation process. The first measure of settlement failure is whether the policyholder filed suit against his insurer, as opposed to settling without suit. The second measure is whether, conditional on the policyholder filing suit, the claim proceeded to trial as opposed to settling out of court.[14] We also use data on the policyholder's settlement demand, as well as claimant- and accident-specific variables, to analyze how the types of claims that proceed to each stage of the negotiation process differ between legal regimes.

Empirical estimates of the effect of bad faith tort liabilty on settlement failure support the predictions of the Priest-Klein model. The strongest finding pertains to the likelihood of an insurance claim to proceed to trial conditional on the case entering litigation. The likelihood of an average claim being closed in trial rose by over 20% among tort states in the earliest years following their regime change, relative to non-tort states. Trial likelihood in tort states thereafter declined steadily, dropping over 10% below other states by the end of the sample. Together, these facts are consistent with theory that selection operates on cases closest to the decision standard. The tort liability regime changed the types of cases selected for trial. It increased the variance of the error with which parties estimated potential trial outcomes. As the decision standard for the tort regime solidified over time, the regime's level of uncertainty changed from a mixture of transitory and permanent levels to a solely permanent level. A narrower range of disputes around the standard survived to trial, so that the density of trials decreased.

We find a similar pattern for the likelihood of litigation, i.e., a policyholder filing suit against his insurer. We estimate a positive effect of bad faith tort liability on the suit probability of an average claim of nearly 5% in the early- and intermediate-phase of bad faith re-

[14]The dataset contains only closed claims. Therefore it does not contain any cases in which the policyholder "dropped" the lawsuit. See Katz (1990) for an empirical study examining drop rates for medical malpractice claims.

form, but a negative effect in the long-term. Regression specifications control for numerous characteristics that exogenously impact the likelihood of settlement, including economic loss claimed, attorney use, degree of injuries suffered, and other state-level legislation on punitive damages and prejudgment interest. Results are robust to sample selection bias with respect to the timing of legal changes, as well as to endogeneity of the time between a claim's filing and settlement. We conclude that the combination of the two changes associated with tort liability—the level of uncertainty and the stakes of the case—created two discrete effects on the likelihood of settlement failure for first-party bad faith auto insurance claims. The short-term effect of uncertainty significantly increased litigation and trial rates among tort states in the 1970s and 1980s. But steadily declining uncertainty about bad faith tort law, in tandem with limits on punitive damages and gradually narrowing doctrinal scope for bad faith tort liability, eventually reduced the scope for disagreement in settlements. This has ultimately reduced the overall strain that settlement failure in first-party auto insurance claims imposes upon the legal system.

While these results are in line with divergent expectations models, they are inconsistent with the full implications of asymmetric information models of settlement. To better distinguish between the two types of canonical models, we examine the economic loss that the policyholder claimed. This constitutes a direct proxy for her settlement demand—typically unobservable in bargaining settings—which is the central lever that drives comparative statics results in asymmetric information models. We find that policyholders' average settlement demand changed non-monotonically through the sample period. This is at odds with both screening-based models of settlement, which predict a monotonic increase, and signaling-based models, which predict a decrease. There is, however, evidence that the settlement demand was correlated with the likelihood of litigation, as both screening and signaling models predict.

This paper makes two primary contributions. The first is to empirically test the divergent expectations model of pre-trial settlement bargaining.[15] We follow Priest (1987), Hughes and Snyder (1995) Waldfogel (1995), Siegelman and Donohue (1995), Kessler, Meites and Miller (1996), Waldfogel (1998), Siegelman and Waldfogel (1999) and Eisenberg and Farber (2003) in testing the Priest-Klein model of settlement.

In contrast to studies that focus on a limiting prediction of the model—50% win rates for plaintiffs and defendants—we assess its broader implication that the selection of disputes for litigation or trial is based on parties' subjective assessments of trial outcomes. Our study is closest to Priest (1987) in estimating the effects of legal uncertainty about a decision standard. We show that litigation and trial rates for a similar set of disputes increased dramatically under a new legal regime, but eventually decreased following the regime's maturation. Combined with descriptive evidence of different accident characteristics between the two regimes, this evidence is consistent with variation in the precision of future trial outcomes driving parties' settlement decisions. More generally, it is also consistent with disputes for litigation or trial being systematically selected on the basis of proximity to a legal standard.

Other empirical evidence on litigation and trial outcomes is also consistent with divergent expectations models. While there is limited evidence for the predictions of asymmetric information models, the literature does show that the direction of comparative statics may depend on the types of disputes considered.[16] Viscusi (1988) and Perloff, Rubinfeld and Ruud (1996) specify models of litigation and settlement in products liability and antitrust claims, respectively. Consistent with both types of canonical models, they find that trial probability increases with the size of the case and decreases with the variance of outcomes and the magnitude of risk aversion. Fournier and Zuehlke (1989), in contrast, study a

[15] A developing line of research documents and explains the existence of partial settlements between the two extremes of completely contingent settlement out of court and completely incontingent litigation in court. See Prescott, Spier and Yoon (2014), Prescott and Spier (2015*a*) and Prescott and Spier (2015*b*).

[16] See Farber and White (1991) and Farber and White (1994) for empirical evidence of screening in medical malpractice litigation.

broader database of civil suits. They find that larger expected payments are associated with greater rates of settlement, an outcome of divergent expectations but not screening with asymmetric information. Johnston and Waldfogel (2002) exploit federal litigation data that identifies attorney representation to show that increased interactions speed dispute resolution and raise the probability of settlement. The present study emphasizes that the variance in parties' subjective probability estimates may change over time to greatly shape the direction of comparative statics.

The second contribution is to estimate the shadow effect of a change in tort liability, specifically, on the likelihood of litigation. We specify a model of insurance claims in the spirit of previous empirical analyses of malpractice claims (Danzon and Lillard, 1983; Hughes and Snyder, 1989), auto insurance claims (Kessler, 1996; Browne and Schmit, 2008), and civil litigation cases more generally (Eaton, Mustard and Talarico, 2005).

Our paper is unique among this group in exploiting variation on the extensive margin of tort law—from no cause of action to a tort-based cause of action—rather than on the intensive margin, from an unrestricted tort right to a tort right restricted by caps on various damages. Perhaps for this reason, many of these studies fail to find significant effects of damage changes on litigation outcomes. Kessler and McClellan (2005) find that tort caps reduce lawsuit incidence by only about 2%. Kessler (1996) finds that tort reforms mandating prejudgment interest for winning plaintiffs increase settlement delay in auto insurance claims, and Eaton, Mustard and Talarico (2005) find no evidence of punitive damages on settlement rates or amounts in civil cases throughout the state of Georgia.

The outline of the paper is as follows. Section 2.2 develops a model of pre-trial negotiation pertaining to bad faith insurance settlements and to the structure of the dataset presented. Section 2.3 details the data used for empirical analysis and presents descriptive evidence consistent with the model. Section 2.4 conducts econometrically estimates the relationship between tort liability and settlement failure over time. We conclude in section 2.5 by discussing implications for the economic theory of settlement.

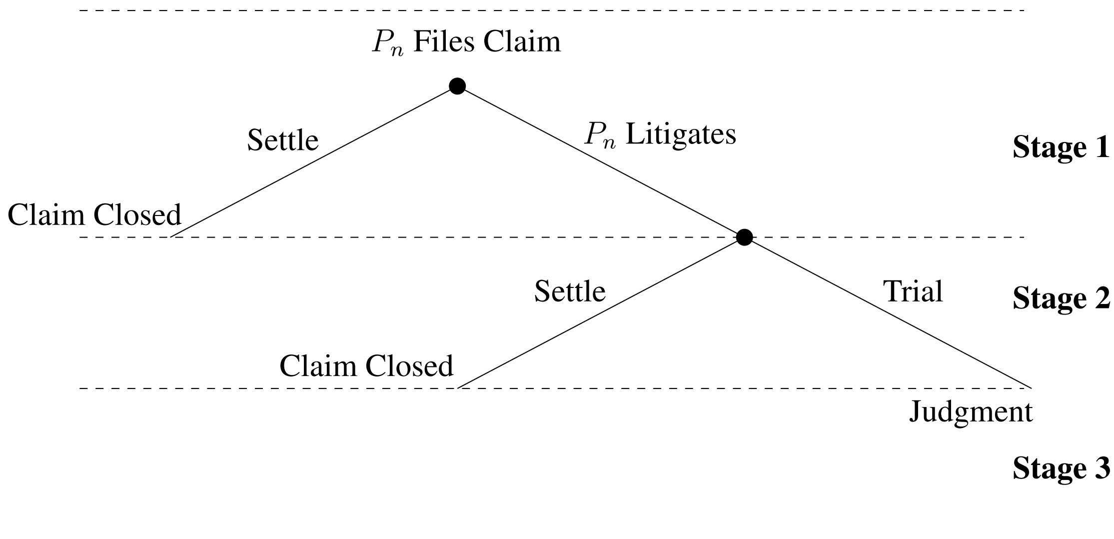

Figure 2.1: The bargaining process

2.2 Theoretical Framework

2.2.1 Model Preliminaries

We adopt a divergent expectations model of bargaining in the form of Priest and Klein (1984) and apply it to the bad faith insurance setting. The model compares the probability of an insurance claim from two different bad faith regimes settling over different time periods. Consider one insurance company, D, that is charged with settling first-party insurance claims from the set of its policyholders, P. Each realized claim x is filed by policyholder P_n, $n \in \{1, \ldots, N\}$. P_n and D bargain over the amount of compensation for the claim.

Settlement occurs if the minimum amount demanded by P_n is less than or equal to the maximum settlement offer from D. There are three discrete stages of the settlement process, indexed by s. Parties hold full information about the timing of the process and the options to settle or prolong the negotiation at each stage. Figure 2.1 represents the three stages of negotiatiation, which culminate in a court or jury trial.

At the first stage, the two parties bargain after a claim has been filed. If they agree on a settlement amount, claim x "closes" within this stage and never reappears. Should they fail to reach an agreement, P_n files suit against D for a breach of insurance contract. P_n

can also allege a supplementary charge that I negotiated in bad faith with P_n.

If P_n files suit, then the bargaining process repeats at the second stage. If they fail to reach an agreement at the second stage, the negotiation over x proceeds to the third and final stage: trial. Parties hold full information about the amount at stake in the trial, judgement $J \geq 0$. This judgment may include compensatory damages for bad faith, if it is alleged. The judgment award J from D to P_n closes the claim in stage three.

If x is in stage one or stage two, P_n and D formulate settlement demands and settlement offers, respectively. They formulate subjective probabilities $\pi^i \in [0, 1]$, $i \in \{P_n, D\}$, that P_n would win trial against D. Party i's optimism about trial success is weighed against his cost of prolonging the negotiation, $r_s^i(J) > 0$, $s \in \{1, 2\}$, where s denotes the stage of negotiation. Let $r_1^i(J) < r_2^i(J)\ \forall\ J$: the cost of settlement failure is strictly greater in the second stage, when trial results, than in the first stage, when lawsuit merely results.[17] Assume also that $\frac{\partial r_s^i}{\partial J} > 0$: litigation cost is increasing in the size of the judgment.

Denoting $r_s(J) = r_s^P(J) + r_s^D(J)$, a first or second stage claim fails to settle in that period if and only if:

$$J\pi^P - r_s^P(J) > J\pi^D + r_s^D(J) \tag{2.1}$$

$$\pi^P - \pi^D > \frac{r_s(J)}{J} \tag{2.2}$$

The standard condition in 2.2 implies that the minimum asking price of P_n exceeds the maximum offer of D at stage s. The assumption that $r_1^i(J) < r_2^i(J)$ guarantees that parties are more willing to settle at the second stage than the first stage.

The defining feature of the Priest-Klein model is its endogenous formulation of subjective trial assessments π^i. Settlement fails to occur if parties hold sufficiently optimistic expectations of their chances at trial. A difference in optimism is determined by two fac-

[17] Cost is comprised of attorneys' fees and opportunity cost of not settling earlier. Attorneys' fees from filing suit include initial fact finding, documentation, case review, and communication between parties. Additional costs for trial include preliminary injunctions, wiitness preparation, and testimony presentation; expenses increase sharply as a case proceeds to trial.

tors. The first is the relationship between the legal facts of the claim and the correctly administered decision standard. Each claim x contains a different mixture of bargaining characteristics. These include the timing and nature of initial contact between policyholder and insurer, the manner of claim appraisal and evaluation, the settlement offer, threats made to agree to the settlement offer, and more generally, all dealings that fall under the umbrella of the insurer's "good faith duty to settle." $Y \in \mathbb{R}$ maps the mixture of characteristics associated with each claim onto a continuous measure of insurer fault. Y^* is the degree of fault necessary for a court to hold the insurer liable for bad faith under the given standard, and Y' is the particular degree of fault that x displays.

The second factor that determines π^i is error with the estimation of Y'. The facts of a claim negotiation map onto an objective assessment of the insurer's wrongdoing, but the application of those facts to the case requires careful consideration and analysis. Parties may be uncertain or disagree over the steps taken after the claim was filed or the appropriateness of various negotiation tactics. They may also disagree over the interpretation of the standard or, particularly in the bad faith setting, the expectation over the future evolution of the standard as it pertains to the claim. This measurement error creates subjective assessments of the insurer's fault. Priest-Klein assume that errors are independently and identically normally distributed, so that draws in opposite directions create optimism bias.

Formally, let $\hat{Y}^i = Y' + \epsilon,\ \epsilon \sim N(0, \sigma^2)$. The probability of settlement failure at stage $s \leq 2$ is:

$$Pr[Fail_s] = F^P\left(\hat{Y}^P\right) - F^D\left(\hat{Y}^D\right) > \frac{r_s\left(J\right)}{J} \quad , \ F^i\left(\cdot\right) = \int_{-\infty}^{\hat{Y}^i} \phi(\epsilon) d\epsilon \tag{2.3}$$

It is important to note that condition 2.3 represents unconditional probabilities of settlement failure. As such, $Pr[Fail_s]$ denotes the proportion of claims from an initial set (each beginning at the first stage) that failed to settle by stage s. Because $r_1^i(J) < r_2^i(J)$, $Pr[Fail_2] < Pr[Fail_1]$: strictly more claims will settle by stage two than by stage one.

Equivalently, $Pr[Fail_2|Fail_1] = \frac{Pr[Fail_1]-Pr[Fail_2]}{Pr[Fail_1]} < 1$: some disputes that failed at stage one will settle at stage two, because litigation costs rise by enough to offset optimism bias.

Condition 2.3 also implies the central prediction of the Priest-Klein model. When all else is equal, claims with Y' close to Y^* are less likely to settle than claims with Y' far from Y^*. If $Y' \approx Y^*$, even a small divergence in error terms results in a relatively large area under the density curve. This enhances the probability of disagreement.

2.2.2 Tort Liability for Bad Faith

In this section we add a stylized representation of the bad faith legal environment to the model. We use the framework to generate predictions of the effect of bad faith tort law on the probability of settlement failure for auto insurance claims at different points in time, as well as to derive conditions under which the empirical setting identifies this effect.

Insurer D operates in two states, A and B, in each of infinitely many discrete periods of time t. D's policyholders P file a total of X_t^l new claims, $l \in \{A, B\}$, in each period t. Claims are filed according to the same distribution with respect to the liability standard, Y, in both states. Specifically, $X_t^l \sim Poisson(\lambda)$.[18] Each stage of the settlement process corresponds to one period t. The framework outlined in section 2.2.1 implies that the X_t^l claims can be partitioned into the subsets $X_t^{l1}, X_t^{l2}, X_t^{l3}$, $\cup X_t^{li} = X_t^l$. With probability $1 - Pr[Fail_1]$, $X_t^{l1} \subset X_{lt}$ will settle in t. With probability $(Pr[Fail_1] - Pr[Fail_2])$, $X_t^{l2} \subset X_t^{l1}$ will settle in $t+1$; and with $Pr[Fail_2]$, $X_t^{l3} \subset X_t^{l2}$ will settle in $t+2$.

There are two possible bad faith regimes: no private cause of action for bad faith, which renders it part of contract law, or tort-based cause of action. States A and B do not recognize a private cause of action for bad faith in periods $[1, \ldots, t-1]$. In this interval, subjective trial assessments are made from the same decision standard, Y^*, in both states.

[18]This assumption can be relaxed in two ways without changing the results that follow. The mean number of claims may differ across states. The mean number of claims may also change over time, as long the change is identical across states. These points are explained further in the empirical identification discussion in section 2.4.

The probabilty of settlement failure across all three stages is equal for both states. At the start of period t, state B applies a tort remedy for bad faith. All claims that are filed in period t are subject to the state's prevailing legal regime, but claims filed before t that are still in negotiation at t remain subject to the law when the claim originated.

Tort adjudication differs in two ways from conventional contract adjudication. First, tort law introduces the possibility of economic damages above the insurance policy's limits, noneconomic damages such as emotional distress, and punitive damages at the discretion of the court.[19] The model therefore specifies that the effect of tort liability is to raise the stakes of a case relative to contract liability.

Second, the change from one body of law to another alters the decision standard Y^*. A new Y^* changes the variance of parties' error, σ^2, because it carries a new level of uncertainty. Priest-Klein allows this uncertainty to change over time. Priest (1987) writes:

> One would expect virtually all changes in legal rules to upset the parties' expectations of outcomes over some period of time, however short. That is, holding constant the permanent level of rule uncertainty, one would imagine that the parties will face greater uncertainty in predictions of the initial application of a new legal standard than they will face in predictions of subsequent applications of the standard. Thus, the transitory level of uncertainty should surely decline over time as study of the rule or actual experience with implementation of the rule accumulates.

A tort change at period t implies that failure probabilities are no longer invariant across state and time. The model outlined in this section, as well as condition 2.3, can be used to derive comparative statics and dynamics of settlement failure with respect to each of the variables that tort law changes. The following two propositions establish testable predictions of this legal change.

[19]Some states have limitations on the extent of punitive damages desipte adjudicating bad faith through tort. Date on punitive damages reforms is described in section 2.3 and used in the empirical analysis.

Proposition 2.1 *Denote the set of claims that close in period t, in state l, and at stage s as Z_t^{ls}. Suppose that the distribution of claims with respect to Y is identical between states A and B in periods $[t-2, t-1, t]$.*

If tort liability raises the stakes of a claim from J_1 to $J_2 > J_1$, then:

$$(J_2 - J_1)\left[\pi^P - \pi^D\right] > r_s(J_2) - r_s(J_1) \iff \frac{Z_t^{Bs}}{Z_t^{B1}} > \frac{Z_t^{As}}{Z_t^{A1}}$$

where $s \in \{1, 2\}$

Proof 2.1 *See Appendix B.1.1.*

Proposition 2.1 represents a familiar result in the law and economics of pre-trial settlement. Cross-sectionally, raising the stakes of a trial J results in two distinct and opposing effects. First, it magnifies existing optimism bias between plaintiffs. This increases the scope for disagreement. Second, however, it increases negotiation effort before litigation or trial. When a case is more valuable, insurers increase their investment into discovering the facts of the policyholder's case by consulting police and hospital reports, conducting independent medical examinations, and appraising the extent of property damage. Plaintiffs also have the incentive to increase spending into discovery and independent evaluation. Increases in the costs of prolonging a dispute reduce the likelihood of disagreement by enhancing the settlement zone.[20] The net effect in period t is therefore ambiguous, and the magnitude of any change is an empirical question.

[20]Note that if parties are risk averse, the same argument applies: the likelihood of settlement failure decreases. The magnitude or impact of risk aversion is not tested, but is mentioned in section 2.4 where pertinent.

Proposition 2.2 *Assume the same conditions as in 2.1. If lower uncertainty reduces the variance of* ϵ *to* $\sigma_{t+1}^{2B} < \sigma_t^{2B}$*, then:*

(i) $\frac{Z_t^{Bs}}{Z_t^{B1}} > \frac{Z_{t+1}^{Bs}}{Z_{t+1}^{B1}}$*, where* $s \in \{1, 2\}$.

(ii) *(a)* $(r_s(J_2) - r_s(J_1)) < J^{**} \Rightarrow \frac{Z_{t+1}^{Bs}}{Z_{t+1}^{B1}} > \frac{Z_{t+1}^{As}}{Z_{t+1}^{A1}}$

(b) $(r_s(J_2) - r_s(J_1)) > J^{**} \Rightarrow \frac{Z_{t+1}^{Bs}}{Z_{t+1}^{B1}} < \frac{Z_{t+1}^{As}}{Z_{t+1}^{A1}}$

where J^{**} *is a sufficiently large constant.*

Proof 2.2 *See Appendix B.1.2.*

Proposition 2.2 constitutes a comparative dynamic prediction of changes in settlement failure over time, both within tort states and between tort and non-tort states. Part (i) generates an unambiguous prediction of the effect of legal uncertainty on settlement rates. It follows from the fundamental properties through which Priest-Klein specificies subjective probability formation. More uncertainty creates a larger probability of settlement failure by increasing the error with which parties estimate the standard a given case, Y'. Whatever is the permanent level of uncertainty associated with the standard for tort law, the possibility of transitory uncertainty adds to this value in the short-term. This implies that a lower percentage of claims will settle without litigation or trial in state B in the period of regime change than in the period(s) following regime change.

Part (ii) combines the earlier two results to compare settlement rates at period $t+1$ between states A and B. Unlike (i), a comparison between states depends not only on parties' uncertainty in evaluating potential trial outcomes, but also in their response to the larger stakes of tort trials. Like Proposition 2.1, response can be characterized by the difference in litigation costs between regimes. Increased litigation costs reduce the probability of settlement failure, while increased effective optimism bias and possibly increased permanent uncertainty raise the probability of settlement failure. If the effect of increased litigation costs is large enough to offset the effects of increased effective optimism bias and permanent uncertainty, then closed claims in state B will have lower levels of settlement failure in the long-term than closed claims in state A. And vice versa if optimism bias and uncertainty dominate litigation spending. The dynamic effect of tort change between states, like the static effect, is therefore also an empirical question.

Before testing the predictions of the Priest-Klein model for auto insurance bad faith, it is worth pausing to assess the implications of asymmetric information models to the same question. These models posit that an uninformed plaintiff makes settlemend demands of an informed defendant. In the auto insurance context, the insurer could expect to have dealt with a range of bad faith disputes or to know whether its actions in negotiation constitute negligent or intentional harm. With respect to tort law's increased stakes, screening models such as Bebchuk (1984) imply that the law increases the policyholder's threat position and therefore his settlement demand. Higher settlement demand requires the marginal insurer to possess a weaker case than otherwise. This unambiguously increases the overall percentage of claims that fail to settle.

Strategic models of signaling add an additional layer of complexity to the standard screening framework by allowing parties to bargain sequentially. One party's decision to make or respond to an offer conveys a signal to the other party, which he uses to update his prior assessment of the other's type. Nalebuff (1987) specifies a two-sided model of asymmetry in which the policyholder makes a settlement demand, the plaintiff accepts or rejects, and the policyholder then chooses whether to proceed to trial. The policyholder is uninformed about the the strength of his case, and the insurer is uninformed about the policyholder's willingness to extend the dispute to its next stage.

In this model, the signal that the policyholder's offer conveys to the insurer creates a "credibility constraint." Because of the double-sided asymmetry, the policyholder must demand a high enough settlement to provide the insurer a credible threat to go to trial. Policyholders raise the value of their demands relative to the non-signaling benchmark. The presence of the credibility constraint can reverse the Bebchuk (1984) prediction. If it binds in equilibrium, then the increased stakes of a tort regime relax the constraint. The policyholder no longer inflates his negotiation position, which *decreases* his settlement demand relative to the non-tort case. This leads the insurer to accept a higher proportion of policyholder offers, and the incidence of settlement failure unambiguously decreases.

Neither the screening nor signaling models specify how uninformed parties form their expectations over their adversary's type. As such, they do not generate testable predictions for how the uncertainty induced by a change to tort law would alter the distribution of insurer types that the policyholder perceives. Likewise, because the magnitude of parties' error is left unspecified, asymmetric information models do not predict comparative dynamics such as those arising from the Priest-Klein model.

2.3 Data and Descriptive Evidence

2.3.1 IRC Data

To test the predictions derived in section 2.2, we employ a repeated cross-sectional dataset of auto insurance claims. Specifically, we use three samples of closed, uninsured motorist (UM) claims from surveys conducted by the Insurance Research Council (IRC).[21] UM is a standard coverage that provides policyholders compensation in the event that they suffer bodily injury in an accident where the other driver is at fault, but that driver does not possess liability insurance. It is consistently interpreted by courts as a type of "first-party" insurance: the policyholder files a claim requesting payment from his insurer for his own loss, rather than for the loss of a "third-party."[22] As Browne, Pryor and Puelz (2004)

[21]The IRC is a nonprofit research organization that has conducted numerous decennial surveys of American insurers since its inception in 1977. It is widely regarded as the leading data provider for economic studies of auto insurance claims.

[22]See Corp. and LLP (2008), described further below. Third-party (a type of "liability") insurance provides coverage to a policyholder for a loss the policyholder brings upon a third-party. Bad faith law treats the two types of insurance differently. Expanding liability for third-party insurance has been less legally contentious than its analog, and states have been quicker to recognize tort liability on third-party claims. See Sykes (1994) for a normative economic analysis of third-party bad faith.

and Asmat and Tennyson (2014) describe, UM insurance provides ample scope for the policyholder to negotiate damages with his insurer. UM claims are therefore particularly well suited to assessing settlement likelihood.

The first portion of the dataset consists of the complete set of UM claims from IRC surveys in 1977, 1987, and 1997. Approximately 60 different insurers were surveyed during a 2-week interval in the year the survey was administered.[23] The dataset contains claims across 42 different states in the U.S.[24] It contains broad state-level variation in bad faith regimes and time. We treat the state and year in which the accident occurred as the indicator of bad faith regime. To remove heterogeneity in claims that may affect the chances of dispute in ambiguous ways, we drop claims with outlying features. They include all claims for which claimed loss exceeded the insurance policy limit, another insurer contributed part of the settlement, or a no-fault regime mandated a minimum claiming threshold.

The results presented in section 2.2 depend on claims from both types of states possessing similar distributions and characteristics. The data contains extensive information about the accident, the claim, and the claimant.[25] We exploit in particular information on the date of insurance filing, first payment, and final payment to proxy for the stage of the dispute, s. We construct a measure of "settlement lag" to compare dispute outcomes from claims in different regimes at the same point in their timeline. Settlement lag is defined as the date of final payment from insurer to policyholder less the date of insurance filing.[26] The model specifies that higher settlement lag is associated with higher proportions of lawsuits and trials. We examine the implications of this variable further in section 2.3.3.

[23]Although the share of the private passenger auto insurance market that these insurers comprised is unavailable, the IRC obtains data from most of the largest providers in the country. For example, Allstate, Farmers, GEICO, Liberty Mutual, Prudential, and State Farm have each participated in at least one of the surveys. The IRC does not survey the same set of insurers each decade, and its data does not identify the particular insurer associated with any claim.

[24]Claims from the following eight states are not represented: Hawaii, Louisiana, Maryland, Massachusetts, Montana, Nebraska, New Jersey, and Pennsylvania.

[25]These variables are described further as they appear in the regression specifications in section 2.4.

[26]It could also be defined as the date of first payment less the date of insurance filing; the two measures are nearly identical. We use final payment in order to account for disputes which may have developed after the first payment but before subsequent payments of a structured settlement.

Each claim contains powerful measures of litigation status that we treat as the primary dependent variables of interest. The first is the presence or absence of a filed lawsuit, and the second is a categorical variable for the stage in which a lawsuit ended. The second measure denotes whether a suit ended before trial, during trial, or with judgment from trial. We render the outcome binary by treating a claim as "tried" if a trial begins, and "settled" otherwise.[27] It also contains a binary variable for attorney use during the negotiation.

2.3.2 Bad Faith Regime Data

The second portion of our dataset tracks the evolution of state first-party bad faith regimes. The assessment of state bad faith law over time is based on the dataset compiled and described by Tennyson and Warfel (2010). They use the comprehensive legal sources of case law found in Stempel (2006), Ostrager and Newman (2008), and Corp. and LLP (2008) to classify bad regimes through 2008.[28]

The case law data can be used to classify states by several different types of first-party bad faith remedies: recognition of any private cause of action; tort-based cause of action; punitive damages; or extra-contractual damages. By 1997, the final year of the IRC dataset, all but five states explicitly recognized a private right of action for bad faith.[29] Legal scholarship and economic evidence has argued that the most relevant distinction between regimes is not the presence of a private cause of action, but instead the presence or absence of tort liability.[30] Tort liability typically allows plaintiffs to collect sums for economic harm above the policy limits, non-economic harm such as emotional distress, and punitive damages. Our empirical treatment of bad faith law therefore classifies states as either tort or non-tort regimes at each point in time. This measure is substantively equivalent

[27] Only 8 of the 467 total claims that resulted in lawsuit were settled during trial.

[28] A simplified, cross-sectional version of this timeline was used by Browne, Pryor and Puelz (2004), who analyze auto insurance claims solely from the 1992 IRC survey. The full version was used in Asmat and Tennyson (2014).

[29] The remaining states are Kansas, Maryland, Michigan, Minnesota, Missouri, and New York.

[30] See Asmat and Tennyson (Forthcoming) for a summary of the legal and economic literature on this topic.

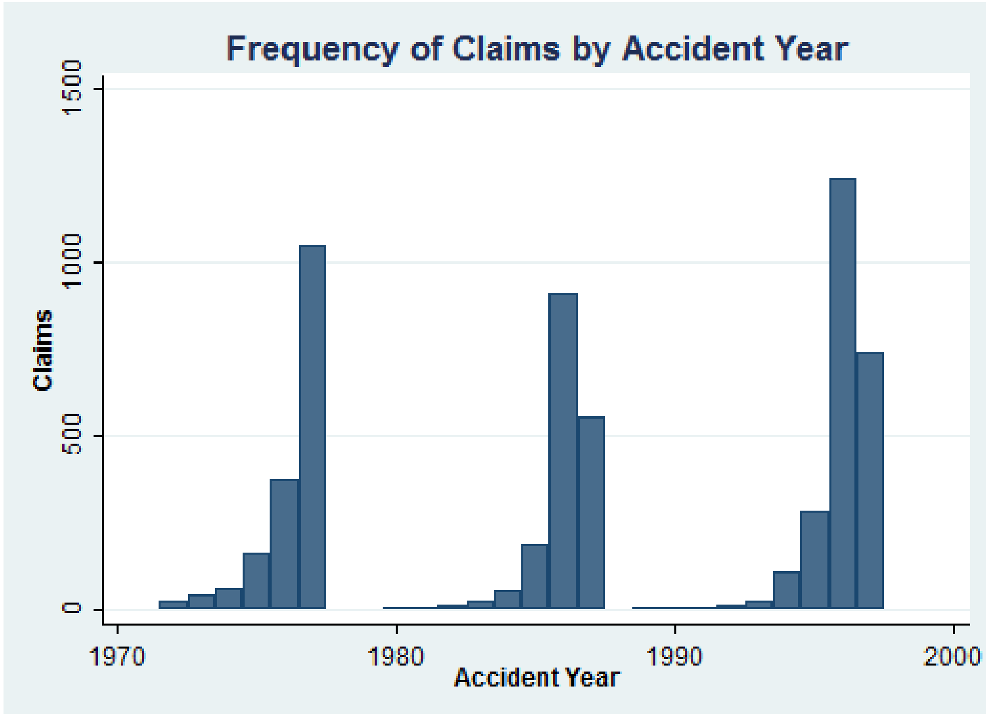

Figure 2.2: Frequency of Claims by Accident Year

to classificaton by the presence or absence of punitive or extra-contractual damages, and results are robust to each of the three classifications.[31]

2.3.3 The Selection Hypothesis

The observations in each IRC survey arise only from claims closed during the two-week interval of sampling. Figure 2.2 displays the number of closed accident claims by survey year for each of the three surveys. Most claims were filed in either the year of or the year preceding the survey, with a small tail of claims filed several years prior. This distribution is relatively constant across the three survey years.

State bad faith reform, on the other hand, occurred non-uniformly over the roughly two decades from the early 1970s to the early 1990s. Table A.1 depicts the final timeline of state bad faith laws with respect to tort liability. It describes whether and when a state ajudicated

[31]Tables of state bad faith laws over time with respect to each of the three remedies are available upon request from the authors.

bad faith through tort law until 1997, the final year of the IRC dataset. States generally either did not enact the tort remedy at any point in their history, or enacted the remedy and retained it through the sample period. The two exceptions are Florida and Georgia, which removed tort liability for bad faith in 1986 and 1989, respectively. To make the legal classification as consistent as possible, we drop claims from these two states.[32]

The interaction of these two parts of the dataset—the structure of the IRC sample and the history of bad faith reform itself—raises the crucial possibility of selection bias. Any claim in an IRC survey for which the accident occurred on or after the state enacted a bad faith tort regime is treated as a tort claim. All other claims are treated as non-tort claims. Simply through random variation in the timing of legal reforms or state populations, tort claims in the sample may be disproportionately from states that switched to tort law several years before an IRC survey year. If most tort claims in the 1977 IRC sample, for example, are from states that switched regimes in 1971-1973 rather than 1976-1977, then there will be a negative correlation between two variables in the dataset. Specifically, as the year that the state switched to tort law approaches 1977, the settlement lag—the time it took to close the claim—would steadily decrease.

Section 2.2 specified a framework in which the settlement lag has a strictly deterministic relationship with the stage of the dispute: prolonging the dispute by one stage prolongs the time period by exactly one stage. The model therefore implies that a greater settlement lag increases the probability of settlement failure (to one, in the limiting case). Moreover, if the year of a state's switch to the tort remedy is negatively correlated with the settlement lag, then it is also negatively correlated with probability of settlement failure. This scenario would generate selection on claims more likely to be have been litigated or tried in tort states. It would result in an upwardly biased estimate of tort law on settlement failure. To the extent that the non-uniformity of regime switches varied from one decade to another—if more states switched to tort law in 1986 and 1987 than in 1976 and 1977, for example—

[32]Estimates with data including claims from these states are available upon request. They do not change the interpretation of any descriptive or regression results.

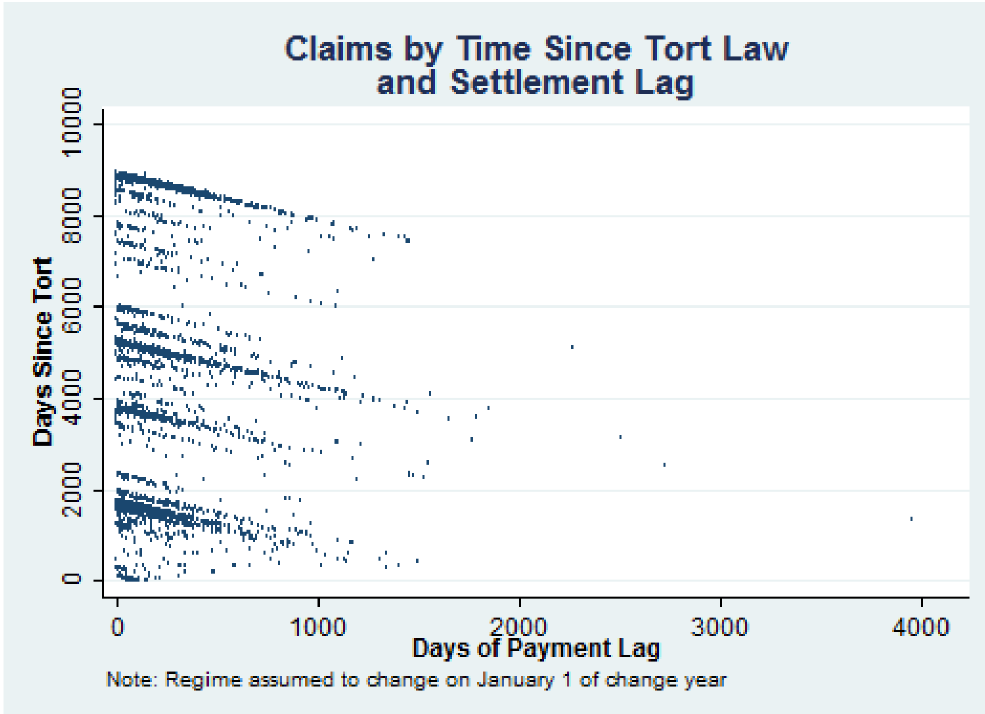

Figure 2.3: Scatterplot: Claims by Regime Switch and Payment Lag

then it would also bias the estimates of regime change over time.

Figure 2.3 addresses this possibility by plotting the settlement lag against the time since a state switched to tort law, in days. The data is encouraging: there is only mild evidence of a negative correlation between the two variables. Throughout the distribution of settlement lag, there is wide variation in the time since switch. We exploit this variation throughout the empirical analysis by conditioning on the settlement lag of a claim. We also account for the possibility that the settlement lag variable is endogenous with bad faith regime. If tort law shifted the distribution of settlement lags leftward, to decrease the chances of a bad faith allegation, then settlement failures would be upwardly biased even conditional on the lag. In section 2.4, we instrument for settlement lag with average settlement lag, by state and month, of a different category of auto insurance claims unaffected by first-party bad faith law.

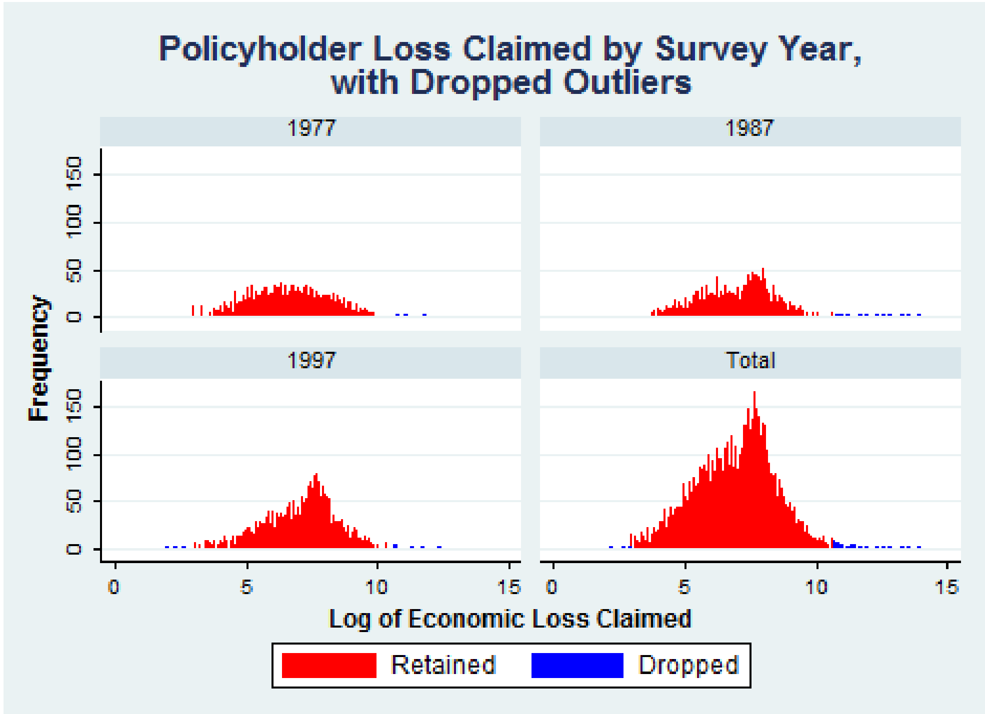

Figure 2.4: Histogram of Economic Loss Claimed, with Outliers

2.3.4 Empirical Evidence

After dropping claims with missing settlement lags, loss claimed, lawsuit or trial status, as well as all other explanatory variables used in regressions, the dataset contains 5,949 observations. Figure 2.4 presents a histogram of the remaining observations. It plots each policyholder's claimed economic loss after the accident, by survey year and for the full sample. The data is approximately symmetric with a rightward skew. To preserve the balance of the sample, we treat economic losses above the 99.9^{th} percentile and below the 0.25^{th} percentile as outliers and drop them.[33] Dropping 59 observations above the 99^{th} percentile and 11 observations below the 0.25^{th} percentile results in a final dataset with 5,879 observations.

[33]Results are robust to the inclusion of these outliers. The six largest losses, all above \$50,000, do alter the magnitudes of claimed loss regressions in the 1987 survey year.

Table 2.1: Settlement Statistics by Legal Regime and Survey Year

	1977			1987			1997		
	Tort	Non-Tort		Tort	Non-Tort		Tort	Non-Tort	
Mean Settlement Lag	211	267	***	242	294	***	280	350	***
	(8.4)	(10.5)		(7.0)	(15.9)		(6.5)	(17.2)	
Proportion Attorney Rep.	0.405	0.427		0.511	0.506		0.445	0.453	
	(0.019)	(0.015)		(0.014)	(0.022)		(0.011)	(0.023)	
Proportion Litigated	0.102	0.114		0.114	0.056	***	0.059	0.093	***
	(0.012)	(0.010)		(0.010)	(0.010)		(0.005)	(0.013)	
N	684	1023		1,216	532		1,943	481	
Proportion Tried | Litigated	0.329	0.077	***	0.072	0.167	*	0.027	0.133	***
	(0.057)	(0.026)		(0.022)	(0.069)		(0.015)	(0.051)	
N	70	117		139	30		113	45	

***$p < 0.01$, **$p < 0.05$, *$p < 0.1$. All values are sample means, with parantheses denoting standard errors. "Loss Claimed" represents the economic losses claimed by the policyholder following accident.

Table 2.1 presents descriptive statistics of pertinent legal variables for the full sample of claims and the subset of claims in which policyholder filed suit. Several results stand out. Consistent with the model, average settlement time is significantly lower in tort relative to non-tort states throughout the sample. This highlights the importance of conditioning on lag to compare litigation outcomes of tort and non-tort states, which we do throughout the paper.

While claimants from both groups of states hired attorneys at a statistically indistinguishable rate, litigation outcomes display remarkably strong patterns over time.[34] The bottom panel shows that 33% of claims in 1977 tort states for which a lawsuit was filed ended in trial. This figure is orders of magnitude higher than other rates of trial both within claim categories in this dataset and across insurance lines in a variety of settings. By 1997, however, the trial rate conditional on suit declined to less than 3%, significantly lower than its 1997 non-tort counterpart. Similarly, the proportion of claims that reach litigation is significantly higher for tort states than non-tort states in 1987. Like trial rates, it also declined to a level significantly lower than non-tort states by the 1997 survey year. These trends are strongly consistent with both parts of Proposition 2.2. Part (i) suggests that transitory uncertainty early in bad faith reform increased the variance of parties' error terms with respect to trial outcomes. Part (ii) suggests that the effect of higher litigation spending dominated higher effective optimisim bias once transitory uncertainty subsided, so that raising the stakes increased settllement in auto insurance.

Figure 2.5 depicts the litigation rates from table 2.1 graphically and nonparametrically controls for variation at the settlement lag-level. The top panel plots the likelihood of an insurance claim being resolved with litigation, by legal regime and survey year. The bottom panel adds a categorical variable for settlement lag. This variable divides claims from each survey into four bins, with each bin corresponding to a threshold for settlement

[34]The claimant is always the policyholder in this dataset, because all claims are first-party. We use the two terms interchangeably.

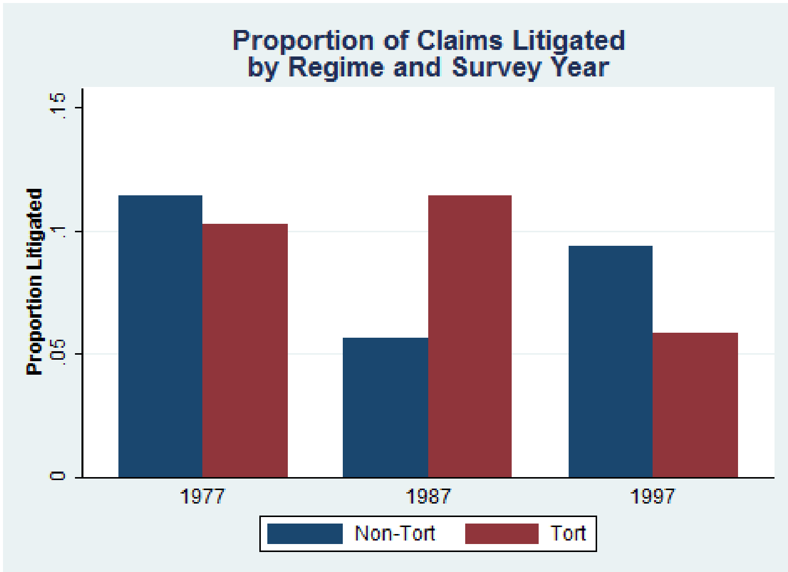

(a)

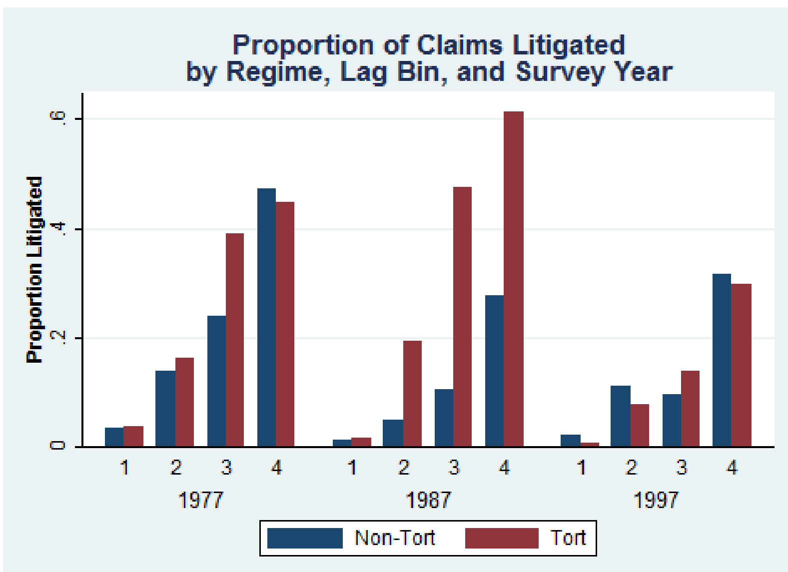

(b)

Figure 2.5: Litigation Probability by Legal Regime and Covariates

lag.[35] Claims from tort and non-tort states are therefore compared at ordered values of settlement lag.

Consistent with the structure of the model, the bottom panel shows clear evidence of data selection by settlement lag: lawsuit proportion rose steadily and markedly as lag increased. Moreover, claims in tort states in the 1977 and 1987 survey years showed higher litigation likelihoods than their counterparts, even after controlling for lag. The effect is particularly strong in the 1987 survey year. The 1997 data, however, shows no evidence of this effect.

Figure 2.6 displays the corresponding bar graphs for trial rates, among the subset of claims to have reached litigation.[36] Unlike lawsuit rates, trial rates conditional on lawsuit did not demonstrate a clear pattern of increase through settlement lag bin, but rates were consistently lower in the first bin than the second.

The graphs reinforce the statistics tabulated in the bottom panel of table 2.1: trial rates rose dramatically in the 1977 survey year. The conditional proportion of claims to trial in tort states exceeded that in non-tort states by at least 15 percentage points in each of the four settlement lag bins. Among the set of claims in the second to fourth bins, well over 30% ended in trial—a startlingly high rate in any civil litigation context. The figure is also noteworthy for a second reason.

It may be hypothesized that a large portion of this steadily decreasing pattern of failure rates was attributable to punitive damages reforms and not waning transitory uncertainty. The history of bad faith law, and civil liability more generally, has seen punitive damage reforms and common law evolution gradually restricting the damages enforcable through tort law. The stakes of tort cases may gradually have been converging to those of contract cases throughout the sample period.

[35]Thresholds for the four bins shown in the figure are 300 days, 450 days, and 600 days, respectively. These values were chosen to highlight variation in the dependent variable; thresholds that correspond to 25% of claims in each bin in each survey year produce similar results.

[36]Thresholds for the four bins are 400, 600, and 900 days, respectively. They are greater than the thresholds in fig. 2.5 because the subset of litigated claims has higher average lags than the dataset at large.

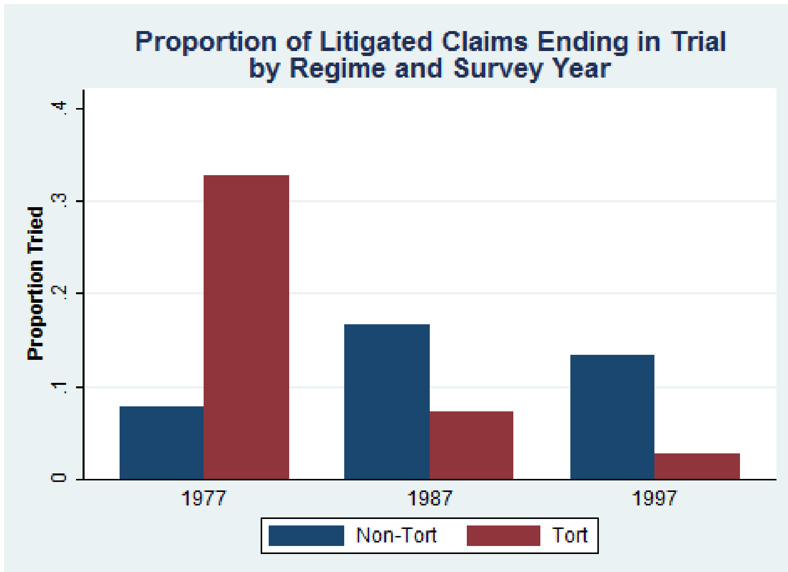

(a)

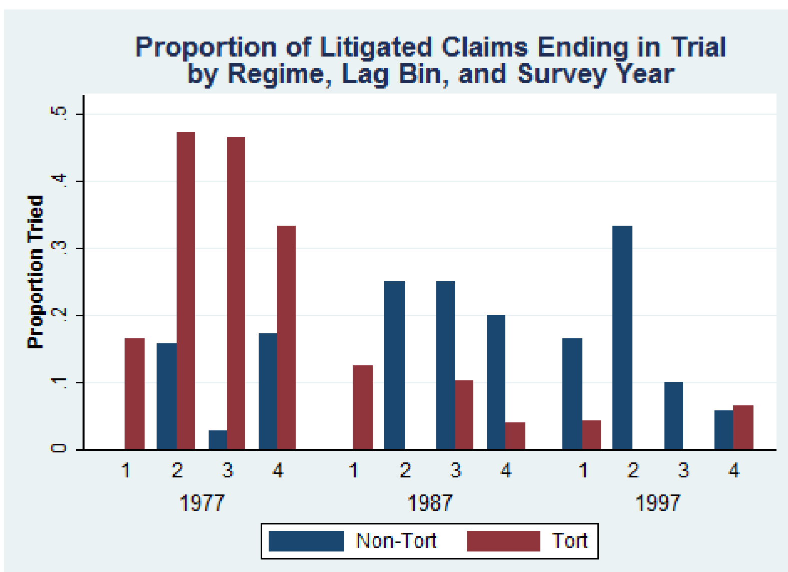

(b)

Figure 2.6: Trial Probability by Legal Regime and Covariates

Figure 2.6 highlights that optimism bias alone cannot explain the trend, however. The reason is that claims from tort states were significantly *less* likely to go to trial or face litigation by the end of the sample, when conditional trial rates dropped under 3%. If optimism bias were solely responsible for reducing the high trial rates of the 1977 survey year, then trial rates in tort rates should be expected to remain at or above those of non-tort states.[37] Instead, this result is consistent with increased litigation spending dominating the impact of optimism bias on net.[38] This effect is only clear in the long-term, when transitory uncertainty had made way for permanent uncertainty.[39] We develop this intuition further in the next section by formulating a full regression specification that controls for punitive damages limits using state-level data.

Finally, we describe other characteristics of litigated claims to show further evidence consistent with the model. Table 2.2 presents descriptive statistics for the subset of cases in which the policyholder filed suit against his insurer, by bad faith regime and survey year. The upper half of the table shows that the claims selected for litigation were different along several dimensions in tort states in the early part of the sample. Claimants who filed suit in 1977 claimed significantly lower economic losses, were significantly less likely to have visited a hospital, and were significantly more likely to have claimed a strain injury than their counterparts in non-tort states.[40] Each of these variables has a clear interepretation for the negotiation of auto insurance claims. Greater claimed losses and injuries require larger payouts from the insurer to settle, on average. Strain injuries are particularly contentious in the insurance setting due to their difficulty in verifiability, and insurers often pursue additional investigation to gauge the extent of the injury.[41]

[37] Furthermore, there are additional damages above the insurance policy's limits that states do not restrict through statute.

[38] Risk aversion is not modeled explicitly, but it could also function identically to increased litigation costs.

[39] As appendix B.1.2 shows, the possibility of tort states possessing lower permanent levels of uncertainty than other states could also produce this effect. It is unlikely, however, because of the well-known complexities in assessing the defendant's level of negligence or intentional harm, as well as his degree of deliberation.

[40] Note that claimants filed significantly lower economic losses, on average, in tort states throughout the sample. Section 2.4 explores this point further.

[41] See Dionne and St-Michel (1991).

Table 2.2: Descriptive Statistics of Litigated Claims
by Legal Regime and Survey Year

	1977			1987			1997	
	Tort	Non-Tort		Tort	Non-Tort		Tort	Non-Tort
Loss Claimed ($)	3,375	4,823		4,357	5,471		5,800	6,059
	(480)	(579)		(443)	(1,376)		(643)	(1,098)
Hospital Stay†	0.414	0.632	***	0.475	0.667	*	0.575	0.689
	(0.059)	(0.045)		(0.043)	(0.088)		(0.047)	(0.070)
Strain Claimed†	0.871	0.692	***	0.863	0.667	***	0.717	0.667
	(0.040)	(0.043)		(0.029)	(0.088)		(0.043)	(0.071)
Age	34.5	28.9	**	33.2	37.0		36.6	32.9
	(2.1)	(1.5)		(1.3)	(3.0)		(1.5)	(2.0)
Employed †	0.343	0.342		0.496	0.567		0.451	0.511
	(0.044)	(0.057)		0.042	(0.092)		(0.047)	(0.075)
Urban Accident†	0.671	0.718		0.583	0.700		0.584	0.689
	(0.057)	(0.042)		(0.041)	(0.085)		(0.047)	(0.070)
N	70	117		139	30		113	45

$***p < 0.01$, $**p < 0.05$, $*p < 0.1$. All values are sample means, with parantheses denoting standard errors. "Loss Claimed" represents the economic losses claimed by the policyholder following accident. All values are in CPI-indexed 1987 US dollars.
† Indicator variable equal to one if condition satisfied and zero otherwise.
Hospital stay equals one if claimant visited ER, overnight, one week, or more than one week.
Strain claimed equals one if claimant reported a neck, back, or other strain.
Employment status equals one if claimant was employed full-time at time of accident, and zero if claimant identified as part-time, student, unemployed, minor, or retiree.
Urban accident equals one if claimant identified accident location as "big city" or "medium city," and zero if identified as "suburban" or "small town."

Furthermore, the table shows that the mean values of these variables in tort states gradually shifted toward equivalence with other states over the course of the sample. By 1997, all three variables are statistically indistinguishable between the two sets of states. The two findings are consistent with the selection mechanism posited by the Priest-Klein model. As the variance in subjective probabilty estimates between the two parties increased in the 1970s, the selection of cases that failed to settle also changed. These cases were not a random sample of the underlying claims filed: they were cases systematically closer than average to the break-even point of the decision standard.[42] Moreover, as the variance in probability estimates decreased over time, the distribution of cases for litigation changed along with it. Sample means of all three variables in tort states steadily converged to the sample means of their counterparts in the 1987 and 1997 survey years.

2.4 Econometric Analysis

2.4.1 Settlement Failure Likelihood

In this section, we test the model's predictions of settlement failure with a sequence of discrete choice regressions. Regression estimation allows us to test a more complete set of predictions from the Priest-Klein model, and to obtain estimated coefficients of the percentage change in settlement failure for claims in the dataset at different sample points. We implement probit regressions of the form:

$$Y_{ilt} = \beta_1 Tort_{lt} + \beta_2 Z_t + \beta_3 Tort_{lt} \times Z_t + \alpha_4 W_{lt} + \alpha_5 X_{ilt} + \epsilon_{ilt} \tag{2.4}$$

Y_{ilt} is defined in two ways, depending on the stage of settlement failure being estimated. The first is whether the claim closed in a lawsuit, and the second is whether the claim closed

[42]They were not a selected sample in all case dimensions, however. The bottom half of the table shows that litigated claims were statistically indistinguishable in age, employment status, and accident location throughout the sample.

with trial, among the subset of cases to have filed a lawsuit. The claim is indexed by *i*, the accident state by *l*, and the accident year by *t*. $Tort_{ilt}$ is coded as one if the state permitted tort law for bad faith in the accident year, and zero otherwise.[43] β_2 captures the changes in parties' settlement behavior attributable to legal uncertainty. Its parameter, Z_t, takes on different values for different regression specifications, as described below. The interaction $Tort_{ilt} \times Z_t$ is the key explanatory variable of interest. β_3 estimates the differential effect of tort law at various points in time.

Identification of β_1, β_2, and β_3 relies on the assumptions outlined in section 2.2. The model restricts the likelihood of settlement failure at stage one or two to be equal across states, at every time period, conditional on bad faith regime. Specifically, it posits that the position of a claim with respect to the decision standard, Y, is equal across regimes. It also assumes that the variance of the error term associated with parties' estimates of Y' is equal within a regime.[44]

These identification assumptions are addressed by conditioning on additional claimant- and accident-level characteristics with the vector X_{ist}. Claim characteristics include the economic loss claimed, final settlement amount, specified injuries and their severity.[45] Accident information includes the presence of hospital stay and duration of any visit (categorical variables for emergency room, overnight, one week, or longer), the size of the city or town in which the accident occurred, and the number of vehicles involved in the accident. Demographic variables include age, gender, marital and employment status.[46]

Identification also requires controlling for ways in which states differ in legal policies other than the bad faith regime. W_{lt} contains two binary variables for state-level laws, puni-

[43]We assume that the dates from table B.1 apply to January 1 of the accident year. Results are invariant to the date chosen.

[44]The variance is not restricted to be constant across time unless it is correlated with bad faith regime. Such changes are picked up through the national-level trend in settlement failure.

[45]The amount of final settlement is also included, but does not distinguish between the insurer's final offer and the final judgment amount in case of trial. See Asmat and Tennyson (2014) for an analysis of bad faith liabilty on settlement amounts.

[46]As described in Section 2.3.1, we also preemptively dropped claims for which claimed loss exceeded the insurance policy limit, another insurer contributed part of the settlement, or a no-fault regime mandated a minimum claiming threshold.

tive damages caps and prejudgment interest, that also constitute deeper tests of the model.[47] Punitive damages caps decrease the stakes of a case by eliminating the possibility of large, outsized judgments, and prejudgment interest increases them by raising the amount owed to a victorious plaintiff.

It is important to account for the possibility that settlement lag is endogenous in the estimation of eq. (2.4). This occurs if tort liability for bad faith induces insurers to speed their negotiation with policyholders in tort states in order to minimize the possibility of litigation. It is instrumented with the average settlement lag in the accident state and year among claims from a parallel IRC dataset on Bodily Injuries. Bodily Injury auto claims are filed against the at-fault driver in accident, rendering them subject to third-party (rather than first-party) bad faith liability. Lags in Bodily Injury claims capture the institutional elements of delay and are correlated with lags from the UM dataset. We assume that this variable is uncorrelated with the error term, ϵ_{ilt}, conditional on the remaining covariates.

Table 2.3 presents the estimation of eq. (2.4) with lawsuit as the dependent variable. In specifications (1)-(3), indicator variables for the 1987 and 1997 survey years act as the time vector Z_t. Specification (2), which adds the full set of explanatory variables, explains substaintially more of the variation in lawsuits than the specification (1). Specification (3) instruments settlement lag with state- and time-level lags from Bodily Injury claims. It captures more variation in the dependent variable otherwise attributed to the time trend.

The primary coefficients of interest hold the expected sign through most specifications: positive in the 1977 and 1987 years, and negative in the 1997 year. The importance of conditioning by (instrumented) settlement lag is evidenced by the fact that it raises lawsuit probability in the 1977 survey year above zero.[48] Specification (3) indicates that a claim with average characteristics in 1987 was nearly 10% more likely to be litigated if it was in

[47]The data listing these reforms by state were obtained from the Insurance Information Institute.

[48]Unreported regressions also show the same sign when attorney use is included as an endogeneous regressor in 2SLS estimation. It is instrumented with a measure of the legal employees in the accident state and year. This data was compiled from the "legal services" 4-digit SIC code in the County Business Patterns series from the U.S. Census Bureau.

Table 2.3: Probit Regression of Lawsuit Indicator
on Tort Liability and Covariates, 1977-1997

	(1)	(2)	(3)†
I(Tort)	-0.010	-0.012	0.005
	(0.013)	(0.011)	(0.011)
I(87)	-0.059***	-0.060***	-0.060
	(0.016)	(0.014)	(0.014)
I(97)	-0.017	-0.028**	-0.037***
	(0.015)	(0.013)	(0.014)
I(Tort) $\times$ I(87)	0.069***	0..054***	0.046**
	(0.020)	(0.018)	(0.018)
I(Tort) $\times$ I(97)	-0.029	-0.021	-0.027
	(0.019)	(0.017)	(0.017)
Ln Settlement Lag			0.056***
			(0.013)
Ln Loss Claimed		0.040***	0.030***
		(0.003)	(0.011)
Pun Dam Cap		-0.017	-0.006
		(0.012)	(0.013)
Prejmt Interest		0.013	0.013
		(0.009)	(0.009)
N	5,879	5,879	5,879
Psuedo R^2	0.015	0.120	

$***p < 0.01$, $**p < 0.05$, $*p < 0.1$. All coefficients are marginal effects calculated at the sample mean of the explanatory variable. All regressions are estimated by MLE.
The dependent variable is equal to one if the claim settled with a lawsuit filed, and zero otherwise. I(87) and I(97) equal one if the claim is from the 1987 and 1997 survey year, respectively. Loss Claimed represents the economic losses claimed by the policyholder following accident. All values are in CPI-indexed 1987 US dollars. Other controls not shown denote employment status, sex, marital status, accident location, injury measures.
† Instruments the natural log of settlement lag with the average natural log of settlement lag for BI claims in the state and year of the accident.

a bad faith tort state than otherwise.

Coefficients on punitive damages caps and prejudgment interest enter negatively and positively, respectively. While neither of these coefficients reaches significance, the signs are consistent with punitive damages causing a small portion of the decline in lawsuit rates in bad faith tort states over time. Both estimates are consistent with the presence of optimism bias, the fundamental property of divergent expectations models.[49]

Note also that settlement lag carries a positive and significant coefficient in specifications (2) and (3). These results are consistent with the Bebchuk (1984) screening model: higher claimed losses are associated with higher likelihoods of filing suit. Unreported specifications with claimed loss regressed on the same covariates indicate that claimed losses were also related non-monotonically to tort liability over the sample period. These results are illustrated in fig. B.3, which plots the mean total loss claimed for the two groups of states in the three survey years. The results suggest that higher claimed losses in 1987 were associated with higher litigation rates in that survey year, and that relatvely lower claimed losses in 1998 were associated with decreased liltigation rates in that survey year. To the extent that claimed loss is endogenous in table 2.3, specifications (2) and (3) overstate its positive effect on lawsuit likelihood and understate the time trends, which nonetheless remain strongly significant.[50]

Table 2.4 presents the analagous marginal effect estimates for trial likelihood. The dataset is made up of claims from the subset of claims for which a suit was filed. Intereacted coefficients are again consistent with increased settlement failure followed by rapidly reduced failure rates by the end of the sample.[51] A claimed strain injury enters positively and significantly. Loss claimed, punitive damages caps, and prejudgment interest are not statistically significant predictors.

[49]While the magnitudes of these coefficients vary based on the point at which the marginal effect is evaluated, the sign does not. Sample means for punitive damage caps and prejudgment interest are 11% and 41% across the sample, respectively.

[50]The next version of this draft will instrument for loss claimed in these two speecifications.

[51]The time trends alone also capture a significant amount of the variation in trial rates; specification (1) has a Psuedo-R^2 of over 11%.

Table 2.4: Probit Regression of Trial Indicator
on Tort Liability and Covariates, 1977-1997

	(1)	(2)	(3)†
I(Tort)	0.156***	0.129***	0.215**
	(0.038)	(0.037)	(0.086)
I(87)	0.073	0.044	0.030
	(0.051)	(0.048)	(0.062)
I(97)	0.050	0.034	0..016
	(0.046)	(0.047)	(0..067)
I(Tort) $\times$ I(87)	-0.234***	-0.224***	-0.270**
	(0.063)	(0.073)	(0.091)
I(Tort) $\times$ I(97)	-0.286***	-0.273***	-0.356***
	(0.063)	(0.066)	(0.092)
Ln Settlement Lag			0.161
			(0.175)
Ln Loss Claimed		0.030**	0.015
		(0.012)	(0.026)
Strain Claimed		0.074*	0.107**
		(0.038)	(0.052)
Pun Dam Cap		0.004	0.025
		(0.061)	(0.082)
Prejmt Interest		0.020	-0.004
		(0.044)	(0 .061)
N	512	512	512
Psuedo R^2	0.113	0.167	

$***p < 0.01$, $**p < 0.05$, $*p < 0.1$. All coefficients are marginal effects calculated at the sample mean of the explanatory variable. All regressions are estimated by MLE.
The dependent variable is equal to one if the claim settled with a trial, and zero otherwise. I(87) and I(97) equal one if the claim is from the 1987 and 1997 survey year, respectively. Loss Claimed represents the economic losses claimed by the policyholder following accident. All values are in CPI-indexed 1987 US dollars. Other controls not shown denote employment status, sex, and marital status, accident location, and other injures.
† Instruments the natural log of settlement lag with the average natural log of settlement lag for BI claims in the state and year of the accident.

The endogeneous settlement lag specification (3) is noteworthy in producing marginal effects much closer to those suggested by the descriptive bar graphs in section 2.3.4.[52] They indicate that tort liability raised the conditional estimate of a litigated claim ending in trial by a remarkable 20% at the sample mean of the explanatory variables in the 1977 survey year. This estimate declined rapidly in the 1987 and 1997 survey years, dropping *below* -20% in 1997. The large magnitude of these estimates is striking.

To gain a deeper understanding of which states drive the trends described above, we decompose the tort indicator variable based on when states switched to a tort regime. Specification (1) in table 2.5 treats a claim as belonging to the tort regime "by" a given survey year if the claim arose from a state that enacted tort law on or before that survey year. The 1977 indicator variable, for example, tags only claims from those states that enacted the tort remedy by the 1977 survey year. The 1987 indicator variable tags claims from those states, in addition to claims from states switched betwen 1977 and 1987.

Specification (2) decomposes tort reform further. It treats a claim as belonging to the tort regime "in" a survey year only if it is from a state that enacted tort law by that survey year *and* after the end of the previous survey year. The "in" 1987 indicator variable, for example, is active only for claims from states that switched by 1987 and after 1977. The next group of variables interacts these indicators with survey year dummys for subsequent years. The coefficient on the "Tort in 1987 $\times$ I(97)" can be interepreted as the effect of tort law on lawsuit probability in 1997 among the second group of states to switch to tort law.

Both specifications in table 2.5 show that the time trends in lawsuit and trial rates were driven by the first states to switch to tort law: those that enacted reform by 1977. The coefficient for the 1977 tort indicator variable in specification (4) implies that those states raised the conditional likelihood of a trial by 21.7% in the 1977 survey year, but reduced the likelihood by 4.3% and 3.0% in the 1987 and 1997 survey years, respectively. There was no corresponding increase in trial rates for the next group of states to pass the tort

[52]The Wald statistic testing for the exogeneity of settlement lag, however, fails to be rejected ($Pr > \chi^2 = 0.44$).

Table 2.5: Probit Regressions of Lawsuit or Trial
on Tort Liability, Time and Covariates, 1977-1997

	Y_{ilt}=Suit		Y_{ilt}=Trial | Suit	
	(1)	(2)	(1)	(2)
I(87)	-0.018***	-0.034***	-0.096***	.003
	(0.004)	(0.006)	(0.024)	(0.053)
I(97)	-0.020***	-0.009	-0.083	-0.004
	(0.005)	(0.006)	(0.027)	(0.040)
I(Tort *by* 77)	0.008*		0.055**	
	(0.005)		(0.027)	
I(Tort *by* 87)	-0.007*		-0.057**	
	(0.005)		(0.027)	
I(Tort *by* 97)	0.006		-0.046	
	(0.013)		(0.125)	
I(Tort *in* 77)		0.001		0.217***
		(0.007)		(0.072)
I(Tort *in* 87)		0.012		-0.045
		(0.023)		(0.030)
I(Tort *in* 97)‡		-0.005		-0.019
		(0.009)		(0.053)
I(Tort *in* 77) × I(87)†		0.070**		-0.043
		(0.035)		(0.034)
I(Tort *in* 77) × I(97)†		-0.006		-0.030
		(0.009)		(0.031)
I(Tort *in* 87) × I(97)‡		-0.014*		-0.073***
		(0.008)		(0.015)
N	5,879	5,879	512	512

***$p < 0.01$, **$p < 0.05$, *$p < 0.1$. All coefficients are marginal effects calculated at the sample mean of the explanatory variable. All specifications instrument the natural log of settlement lag with the average natural log of settlement lag for BI claims in the state and year of the accident, and estimate by MLE.
† denotes Wald test between coefficients in specification (2) rejected for equality at 1% level.
‡ denotes Wald test between coefficients in specification (4) rejected for equality at 10% level.

remedy, the 1987 group: rates continued to decline among those states as they changed regimes Figure B.2 reinforces these results graphically. It shows that the increases in trial rates among tort states in 1977 were not replicated among those states switching to tort between 1977 and 1987, or between 1987 and 1997.

These findings are also consistent with the evolution of bad faith law on a national level. As Tennyson and Warfel (2010) detail, common law in California and Wisconsin formed two noteworthy and distinct standards for first-party bad faith law 1973 and 1978, respectively. Abraham (1994) relates the consensus among legal scholars of the time that liability for bad faith was expected to continue to expand. In fact, however, California and Wisconsin would establish the two primary models of tort law for other states to follow.[53] The 1970s constituted a more fragile time in the development of common law—and therefore to parties' perceived uncertainty regarding liability—than future years would. Figure B.4 shows that the conditional trial trend is fully robust to removing California from the sample, denoting only other states to pass the tort remedy as treated states.[54]

2.5 Conclusion

This paper has exploited state- and time-level variation in the adjudication of tort law for bad faith insurance law to show evidence consistent with the Priest-Klein model. Data from IRC samples in 1977, 1987, and 1997 shows strongly positive and significant effects of tort liability on settlemetn failure rates up to 20% in the short-term, and similarly negative effects in the long-term. The evidence is consistent with the tort regime increasing the variance of parties' subjective trial estimates and increasing the scope for disagreement in the short-term. It is also consistent with one of the predictions of asymmetric information

[53] See Asmat and Tennyson (Forthcoming) for further discussion.

[54] The sample size of litigated and tried claims drops significantly after removing California, particularly in the 1977 survey year (where fewer states adjudicated bad faith through tort). The next version of this draft will update the dataset to add UIM and other types of claims to the sample, increasing size and geographical variation.

models of settlement: policyholder settlement demands were positively correlated with lawsuit probability throughout the sample period.

Although the data cannot isolate the effect of optimism bias specifically, it is likely that a portion of the stark short-term increase in settlement failure was also attributable to the possibility of increasingly large stakes. The 1970s and 1980s saw a national surge in large, multimillion judgments for plaintiffs across civil liability cases. In the subsequent years, states passed numerous reforms to limit the payouts for general and punitive damages, noneconomic damages, and prejudgment interest. Estimates of punitive damages caps and prejudgment interest policies are consistent with optimism bias explaining a portion of the time trend in settlement failure, but they do not come close to explaining all of the variation.

There are also specific factors in the evolution of first-party bad faith insurance law that may have made the magnitude of results presented in this paper especially large. First, there has been debate about the intent and scope implied by the standard that a majority of tort states have imposed. The standard, based on a 1978 Wisconsin Supreme Court case, is understood to lie between the traditionally clear intentional and negligence benchmarks.[55] As Sykes (1996) explains:

> Despite the efforts of the *Anderson* court to define bad faith breach as an "intentional" tort, even a superficial reading of the opinion reveals elements of a negligence standard...The court and perhaps the jury must make some judgment about the objective "reasonableness" of the insurer's actions.

In practice, it is fairly clear that the general trend of this interpretation was to raise the standard required to show bad faith. Abraham (1994) notes that state courts made a point to limit punitive damages following the early years of bad faith tort remedies in the 1970s. Feinman (2012) argues based on more recent cases that bad faith liability has continued to narrow. All of these factors are consistent wtih the decreasing litigation probability trends shown in this paper. Further research is needed to disentangle their impact from those

[55] *Anderson v. Continental Insurance Co.*, 85 Wis. 2d 675, 271 N.W. 2d 368 (1978).

of tort law more generally. This is especially important with respect to the largest policy finding of this paper: that claims from bad faith tort regimes were less likely than claims from their counterparts to face lawsuit or trial by the end of the sample period.

Notwithstanding these particular features in the history of bad faith law, however, this paper has highlighted the importance of a central element of the Priest-Klein model. Because Priest-Klein formulates plaintiff and defendant trial expectations relative to a decision standard, it places first-order emphasis on the variance in parties' expectation of a decision standard. The model is thus well-suited to predicting the effect of the inherent uncertainty that arises when a new legal right is introduced. The tort remedy for bad faith is hardly unique in this regard: practically all changes in substantive law are accompanied by some change in standard. When evaluating comparative statics of a legal change, it is therefore critical to take into account the uncertainty that both parties in pre-trial negotiation perceive. This point is not fully captured by asymmetric information models of settlement.

CHAPTER 3

Empirical Screens of Collusive Pricing: Evidence from a Korean Petrochemical Cartel

This paper uses knowledge of cartel pricing strategies disclosed during an antitrust investigation into South Korean petrochemicals to generate a screen for detecting collusion in commodities markets. This screen is based on differences in the responsiveness of output price to input costs before and after a reversion from collusion to competition. In particular, it examines output price asymmetry with respect to crude oil input prices. Using prices before and after the cartel end date, it constructs measures of the price-cost margin and input price asymmetry for two specific petrochemicals. styrene and toluene. Coefficients from error correction models demonstrated inverted price asymmetry at 1% confidence levels. The conclusions drawn are as follows: the cartel altered supply behavior to enforce collusion, the breakdown in the agreement resulted in systematic changes in the relationships between price and cost, and error correction estimation measures can be used to screen for collusion of this type.

3.1 Introduction

Since the Antitrust Division of the U.S. Department of Justice enhanced its corporate leniency program in the mid-1990s, it has uncovered a persistent stream of international price fixing cartels.[1] The success of the amnesty program has seen many other jurisdictions around the world follow suit and prosecute additional cartels.[2] These cases make clear that firms in many of the world's largest industries have the ability to reach and sustain collusive equilibria today despite the well-known challenges of doing so.

Several recent cases involve cartels in fuel-intensive industries.[3] These cases are especially noteworthy because of the widespread availability of data on crude oil prices around the world. A significant body of empirical literature in industrial organization has used such data to find that prices rise "like rockets" and fall "like feathers" with respect to crude oil prices. The phenomenon occurs when output price displays more responsiveness to input price increases than to input price decreases.[4] There are both procompetitive (Yang and Ye, 2008; Tappata, 2009; Cabral and Fishman, 2012) and anticompetitive (Rotemberg and Saloner, 1986; Haltiwanger and Harrington, 1991; Lewis, 2011) theories of oligopoly pricing that are consistent with asymmetric pricing.

A separate strand of literature uses high-frequency, readily obtainable price data to empirically screen for the presence of collusion in market settings.[5] Abrantes-Metz et al. (2006) and Bolotova, Connor and Miller (2008) use oligopoly models to formulate hypotheses of equilibrium pricing that are consistent with collusion but inconsistent with

[1] Since 1995, the agency has fined firms from over 100 different cartels for over $9 billion. See http://www.justice.gov/atr/public/criminal/sherman10.html.

[2] In addition to the European Commission, competition authorities in Australia, Brazil, and South Korea, among others, have prosecuted cartels as part of amnesty programs over the last decade.

[3] They include air cargo (http://www.justice.gov/atr/public/press_releases/2008/234435.pdf), passenger fares (http://www.justice.gov/atr/cases/f225500/225524.htm), freight forwarding (http://europa.eu/rapid/press-release_IP-12-314_en.htm?locale=en) and retail gasoline (http://www.competitionbureau.gc.ca/eic/site/cb-bc.nsf/eng/03079.html)

[4] See Borenstein, Cameron and Gilbert (1997), Bachmeier and Griffin (2003), Verlinda (2008), Lewis (2011) and references therein.

[5] A related literature in auctions employs other methods to empirically test for collusion; see Porter and Zona (1993, 1999) and Bajari and Ye (2003).

other types of supply. The screens proposed in these studies estimate the first and second moments of the pricing distribution. They do not estimate screens of output price using input prices, or specify models of the degree of price asymmetry during collusion relative to competition.

This paper tests for screens of collusion based on explicit knowledge of how and when a cartel in a petroleum-based market operated. It extends the scope of the screening literature by specifying a clear hypothesis of how collusion affects the degree of price asymmetry relative to competition. It focuses on the case of two petrochemicals products, styrene monomer and toluene, whose South Korean manufacturers were fined by the Korean Fair Trade Commission (KFTC) in 2008. During its investigation, the KFTC disclosed that cartel participants adhered to product pricing formulae to determine collusive behavior. One of the products, styrene, based price directly as a function of crude oil and other inputs; the other, toluene, based price on lagged toluene prices.

This knowledge is used to generate a simple screen with respect to the price asymmetry between the petrochemical price and the crude oil price. Conditions are shown under which adherence to the styrene price strategy would result in an elevated degree of price asymmetry during collusion. This serves as the basis for descriptive and regression measures that are proposed as possible markers of collusion in fuel-intensive industries and commodities markets.

Empirical analysis employs a dataset of weekly, market-wide Korean spot prices for styrene and crude oil, before and after the periods of alleged cartel behavior. It combines this data with weekly crude oil prices during the same time period. The data is used to provide evidence that firms changed supply behavior on or near the dates specified by the KFTC. The timeline is used to generate descriptive measures of price, time-series variation, Pearson product-moment correlation, and adjustment response for both products. It is also the basis for an error correction model estimating the degree of price asymmetry before and after the break date.

Results show that prices in the Korean styrene and toluene markets moved like "rockets and feathers" before the breakup of the cartel, but like "feathers and rockets" after the end of the conspiracy. Consistent with the formulae, the styrene price demonstrated stronger asymmetry than the toluene price during the cartel period and price-cost inversion after the cartel period. Mean adjustment ratios were estimated to fall by 38% and 15% following cartel breakdown in styrene and toluene, respectively. The corresponding coefficients on positive and negative cost changes from the empirical model demonstrate inverted price asymmetry after the break date, with differences in coefficient magnitudes significant at the 1% level.

Other descriptive measures are also consistent with discrete and significant changes to supply behavior after the cartel breakdown, particularly for styrene. Linear correlation declined to nearly zero for styrene and 67% for toluene in the post-cartel sample period. In contrast, measures of first and second moments of prices and margins were not inidicative of collusive behavior in either product. I conclude that the cartel altered supply behavior to enforce collusion, that the breakdown in this agreement resulted in inverse price asymmetry, and that descriptive and error correction estimation measures can be used to screen for collusion of this type.

In addition to the screening literature, this paper supplements work that exploits the institutional features of a particular "hard core" price fixing cartel to draw empirical implications. Levenstein (1997), Roller and Steen (2006), Genesove and Mullin (2001), Clark and Houde (2013, 2014) study the internal rules and organization of cartels to assess their effectiveness in sustaining incentive compatibility. In this study, in contrast, the structure of the cartel is used to generate predictions of the change in price-cost measures before cartel breakup relative to after breakup. Knowledge of the cartel strategy provides a rigorous theoretical basis for collusion that shares attributes with models of endogeneous antitrust enforcement including Harrington (2004). The hypothesized empirical screen of collusion transfers readily to other fuel-intensive industries.

The remainder of the paper is organized as follows. Section 3.2 describes the pertinent features of the Korean petrochemical market and the details of the cartel investigation. Section 3.3 introduces the data under study and summarizes descriptive relationships that contrast cartel pricing before and after the end date. Section 3.4 estimates empirical models of price-cost asymmetry using vector error correction specifications, and section 3.5 highlights the conclusions for screening methodologies.

3.2 Industry Background and Cartel Investigation

3.2.1 Industry Background

The market for petrochemicals in South Korea displays several characteristics long identified as significant in reaching and maintaining collusive equilibria. The market structure is oligopolistic: most petrochemicals are produced by three to five large firms and a smaller fringe of competitors. Table C.1 displays a snapshot of market shares by firm for Styrene and Toluene during 2007, a select year of data availability. Five unique firms accounted for over 90% of the market for both products in 2007.[6] Several of the largest firms are Korean conglomerates, which own industry-leading market shares in areas ranging from telecommunications, media, electronics, and construction.[7] Many of the largest manufacturers also compete in multiple markets and have pled guilty to explicit price fixing in a variety of industries.[8] Prices are widely available at weekly and monthly intervals from two market research firms, ICIS and Platts. There is also a robust spot market for resale and trade.

Two properties make the petrochemical industry in Korea particularly attractive to study. First, production is based heavily on the availability and price of crude oil, an inter-

[6]Globally, Korea ranks between the fourth and sixth largest petrochemical producers in the world, behind the U.S., Germany, Japan, and Saudi Arabia (Korean Petrochemical Industry Association).

[7]Major firms are also Korean-owned and independent of a controlling stake from foreign multinationals. Two cases of partial (but non-controlling) foreign ownership include GS Caltex (Chevron), and S-Oil (Saudi Aramco).

[8]See Asmat (2015) for an empirical analysis of the memory chip cartel and references therein.

nationally traded commodity with readily accessible historical price data. Second, Korean crude oil price is determined exogenously from the petrochemical price: none of the country's petrochemical firms is involved in the direct resale, stockpiling or shipping of crude oil. Figure C.1 shows that the country currently imports 100% of its crude oil.[9] Integrated refineries purchase crude oil and convert it to the distillate naphtha, which in turn produces the various types of petrochemicals after undergoing a cracker process. Toluene is produced directly from naphtha, while styrene is produced from a combination of naphtha, liquefied natural gas, associated gases and ethylene. In addition to its role as a feedstock, petroleum is also used to provide the energy necessary to fuel the chemical reactions which take place. The process is depicted graphically in fig. C.2.

Each product has particular chemical properties that make it homogeneous from both buyers' and sellers' perspectives. The chemical makeup also renders each product suitable for specific end uses. Styrene monomer is a slightly yellow, hazardous liquid produced from the olefins branch of the petrochemicals family. It is sold downstream primarily to synthetic rubber manufacturers; a smaller group of buyers use it to make unsaturated polyesters.[10] Toluene is a colorless, liquid aromatic petrochemical with properties of both rubbers and plastics. It is sold principally to manufacturers of flexible foam, polyvinyl chloride, and polyester; these producers in turn sell to retailers of mattresses, furniture, and automobile parts.[11] Sales are both domestically and internationally; most exports are made to China.[12] Table C.2 shows that firms exported less than half of their production during the years in and surrounding the cartel period.

[9] The largest upstream oil and gas companies are state-owned, including Korea National Oil Corporation.

[10] Synthetic rubber takes the form of polystyrenes including stryrene-butadiene rubber (SBR) and acrylonitrile butadiene styrene (ABS). Many of these manufacturers have pleaded guilty to separate counts of price fixing around the world; see http://europa.eu/rapid/press-release_IP-08-78_en.htm?locale=en.

[11] Two of these downstream markets have also faced fines for international cartels. See http://europa.eu/rapid/press-release_IP-15-5214_en.htm (flexible foam) and http://www.justice.gov/opa/pr/nine-automobile-parts-manufacturers-and-two-executives-agree-plead-guilty-fixing-prices (auto parts).

[12] Korean's petrochemical industry began "on the assumption that...virtually all its surplus petrochemicals and plastics production could be shipped to [China]" ("Korea's Capacity Fears, *Chemical Week*: 2/19/1997).

3.2.2 KFTC Investigation

In 2008, the KFTC began investigating price-fixing among producers of a range of Korean petrochemicals. The investigation studied here closed in June 2008 and imposed fines of about \$12 million on the largest producers of the two chemicals under study.[13] Certain details about the cartel were publicly documented and made available by the KFTC. Styrene and toluene had domestic market shares of 94.6% and 82.5% at the time of investigation, respectively. The KFTC also disclosed the time period in which firms in each product professed to colluding. The styrene cartel officially lasted from October 2000 to September 2004, while the toluene cartel lasted from January 2002 to June 2005.[14]

The most interesting feature of the KFTC's investigation was the disclosure of pricing formulae to which cartel members in each product agreed during collusion. For example, the formula for styrene was given as follows:

$$P_t = \left[\left(\frac{\bar{C}^1_{m-1} + \bar{C}^2_{m-1}}{2}\right)(0.7)\right] * 1.08$$
$$+ \left[\left(0.804 * Benzene_t + 0.280 * Eth\bar{y}lene_q + Won238/kg\right)(0.3)\right] * 1.08$$

This strategy demonstrates a strong relationship between styrene price in period t and input costs of crude oil, benzene, and ethylene. P_t is the price of styrene at time t; $\bar{C}^i_{m-1}$ is the average price of crude oil as listed in the Platts and ICIS data in the previous month, respectively; $Benzene$ is the international price of benzene, and $Eth\bar{y}lene_q$ is the international quarterly average price of ethylene. Specifically, the following hypothesis can be formulated:

[13]For firm-level fine amounts, see http://www.icis.com/Articles/2008/06/23/9134408/s-korean-firms-fined-12.3m-for-pricing-cartel.html. The same investigation fined producers of three other petrochemicals: monoethylene glycol (MEG), diethylene glycol (DEG), and ethylene oxide (EO). Price data for these chemicals was not available for purchase at the time of data collection from ICIS; see section 3.3 more details.

[14]As is typical in criminal proceedings for "hard core" cartels, the dates were agreed upon after a period of negotiation and represent only possible start and end dates for collusion.

Hypothesis 1 *Assume that firms adhere to the formula outlined above in the cartel period. If the degree of price-cost asymmetry between benzene and crude oil on the one hand, and ethylene and crude oil on the other, is greater than one, then the degree of price-cost asymmetry between styrene and crude oil will increase during the cartel period.*

Hypothesis 1 predicts that the degree of styrene-crude oil asymmetry is magnified during the collusive period, as long as there is a pre-existing asymmetry of the other input prices with respect to crude oil. When the price of crude oil rises, it will be fully passed on to the styrene price at the rate of $(0.7) * (1.08)$. To the extent that crude oil price raises the price of the other inputs, those increases will also be passed on at the specified rate. When the price of crude oil falls, it will decrease the price of styrene, passing on a decrease at the rate of $(0.7) * (1.08)$. But to the extent that benzene and ethylene do not drop as sharply as crude oil, the decline will be stemmed. Thus, the degree of price asymmetry will increase.

The KFTC's investigation revealed a different formula for firms colluding in the price for toluene.

$$P_t = \left(\bar{P}_{m-1}\right)(\alpha * Misc\ Expenses)$$

The strategy for toluene price is based on a function of lagged toluene price and a positive multiple, α, times miscellaneous expenses. This strategy does not base output price on the price of crude oil or any other inputs. It therefore does not provide a prior basis to hypothesize an increase or decrease in the degree of toluene-crude oil asymmetry, and the presence or absence of a change is an empirical question.

3.3 Data and Descriptive Evidence

3.3.1 Data Sources

Empirical analysis utilizes historical market-wide price data from the market research firm ICIS. ICIS is a leading provider of price data and market research for chemical, energy, and fertilizer products around the world. The firm's price collection process is detailed and well-cited. Its region-specific employees consult buyers, sellers, and traders worldwide to independently verify price quotations. Its chemical database includes over 180 products, including most of the world's highly traded petrochemicals and feedstocks. It is a primary source for price negotiations between wholesalers, distributors and purchasers.

Specifically, the data purchased from ICIS for this study is comprised of weekly, Free-on-Board import spot prices for two Korean petrochemicals: Styrene Monomer and Toluene.[15] For each product, the time spans two full calendar years prior to the alleged start of the cartel, through the full duration of the cartel period, and two additional calendar years after the alleged end of the cartel.[16] Prices are quoted at the end of business each Friday, exclude the final two weeks of the year. They consist of a weekly high and low value, reported in US dollars/Tens of Metric Tons.[17] I supplement petrochemical prices with international crude oil prices from the Bloomberg database. Specifically, I use Iran "Light" crude oil priced for export to Asia.[18] The price in US dollars at the close of trading each Friday is reported as the (weekly) unit of observation.

[15]ICIS does not offer export prices for the Korea region.

[16]Additional periods of historic data before or after the cartel period were not purchased.

[17]Prices are non-deflated throughout the analysis because of the sensitivity of output prices to oil prices. The results are robust to deflating both indices by monthly Korean CPI data from KOSTAT, Korea's national statistics agency.

[18]Results are not sensitive to the grade of oil used; covariance between Iran light and several other internationally traded grades averaged over 0.97 during the sample period.

3.3.2 Cartel Break Evidence

This section outlines descriptive evidence to support the contention that firms changed supply behavior during the data sample period. The change in behavior corresponds to the date of cartel breakup alleged in the KFTC investigation, lending credence to the institututional evidence. Figure 3.1 presents graphs of petrochemical price, crude oil price, and the margin. Panel (a) depicts the styrene relationship and panel (b) the toluene relationship. The gray shaded area represents admitted dates of collusive activity.

The figures show that the trend in price-cost margin was gradually increasing during the cartel period of both products. Styrene margin rose until September 2004, the month of cartel breakup, before decreasing for the rest of the sample. The same pattern held for toluene, where the price-cost margin increased from the start of its cartel period in January 2002 through September 2004. Margin thereafter declined sharply through the rest of the conspiracy period and the sample.

Table 3.1 presents summary measures of the first moments, linear correlation coefficient, margin, and adjustment ratio for each of the three periods of supply behavior, as given by the KFTC. Let the petrochemical price in time t be denoted P_t and crude oil price in time t C_t. Following Clark and Houde (2014), the adjustment ratio of price asymmetry is defined as follows.

$$AdjustmentRatio_t = \frac{\sum_{t=1}^{T} \Delta P_t^+}{\sum_{t=1}^{T} \Delta P_t^-} \Bigg/ \frac{\sum_{t=1}^{T} \Delta C_t^+}{\sum_{t=1}^{T} \Delta C_t^-}$$

This ratio summarizes the mean output price increases divided by the mean output price decreases and standardizes by the corresponding measure in the input price. Values of one represent perfectly symmetric pricing; values greater than one indicate asymmetric pricing; and values less than one indicate inverse asymmetric pricing.

The descriptive statistics in Table 3.1 are noteworthy in highlighting dimensions of

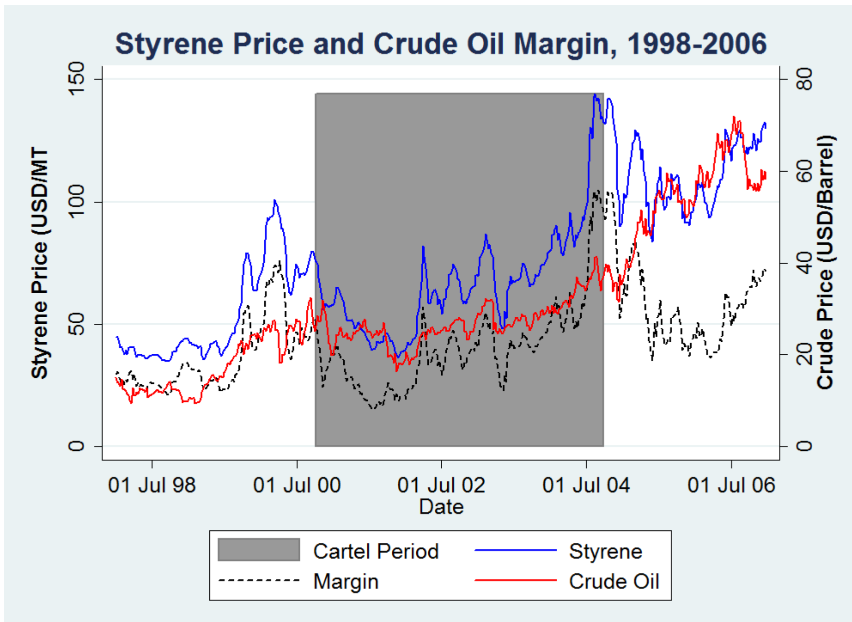

(a) Styrene

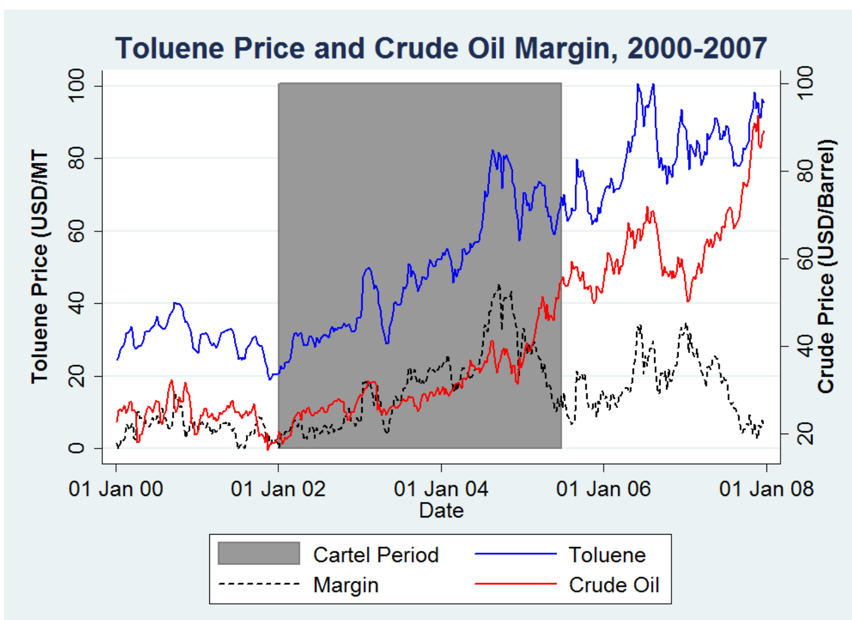

(b) Toluene

Figure 3.1: Chemical-Crude Price Margins

Table 3.1: Summary Statistics by Chemical:
Pre-, During-, and Post-Cartel

	(1)	(2)	(3)
	Pre-Cartel	Cartel	Post-Cartel
Styrene			
Mean	53.82	67.33	112.00
Variance	379.03	494.17	202.77
Corr Coefficient $_t^*$	0.845	0.867	0.029
*Margin**	36.39	40.42	58.049
Adjustment Ratio	1.16	1.38	0.86
N	143	208	116
Time Period	1/98 - 9/00	10/00- 9/04	10/04 - 12/06
Crude Oil			
Mean	17.42	26.91	53.95
Variance	40.04	22.44	94.90
Toluene			
Mean	30.30	48.91	80.91
Variance	21.89	273.77	97.27
Corr $_t^*$	0.758	0.833	0.673
*Margin**	5.94	17.55	17.65
Adjustment Ratio	0.98	1.15	0.98
N	104	181	130
Time Period	1/00-12/01	1/02-6/05	7/05-12/07
Crude Oil			
Mean	24.36	31.37	63.26
Variance	11.10	59.49	92.54

* Denotes Pearson product-moment correlation coefficients. Sample estimates of correlation, margin and adjustment ratio calculated with respect to crude oil price. Price units are in USD/tens of Metric Tons.

variation between hypothesized cartel and non-cartel periods. Two measures of the price-cost relationship, linear correlation coefficient and mean adjustment ratio, were higher for both chemicals during the cartel period than either before or after the cartel period. The linear correlation for styrene declined to near zero after September 2004, consistent with a significant change in pricing behavior.[19] The time-series variation was also markedly higher during the cartel period.[20]

The table also suggests that the largest difference in supply regimes was at the end of collusion rather than at the start, consistent with the collapse of the cartel through a price war. The change from the hypothesized pre-cartel period to the cartel period was not as discrete, consistent with descriptive evidence from Abrantes-Metz et al. (2006), Clark and Houde (2013), and Clark and Houde (2014). The trend is also in line with models of endogeneous antitrust detection posited by Harrington (2005) and Harrington and Chen (2006): price reaches the collusive level gradually rather than suddenly, but falls from the collusive level suddenly rather than gradually.

Figures 3.2 and 3.3 examine this point further by constructing residual plots. The residual plots display diagnostic evidence from OLS and AR(1) regressions, respectively, of petrochemical price on crude oil price. Panel (a) shows clear evidence of autocorrelation throughout the sample, obscuring the inspection of any trend between cartel periods. Partial autocorrelation (PAC) tests displayed in fig. C.3 and fig. C.4 show evidence that each process is approximately integrated of order one. This implies that AR(1) regressions are sufficient to observe residual time trends.

Panel (b) depicts the residuals from the AR(1) model. The results show a much cleaner relationship, without any obvious autocorrelation over the sample period. They are also consistent with the descriptive statistics: the variance of the residuals changes markedy at

[19]Undisplayed cross-correllellogram figures also show significantly altered pricing in the post-cartel period for styrene. In particular, the correlation between styrene in period t and toluene in periods $(t-1, t-2, ...t-20)$ displays a negative and increasing relationship. This stands in contrast to positive cross-correlations at every other point in the sample.

[20]This finding echoes the analysis of the lysine cartel in Bolotova, Connor and Miller (2008).

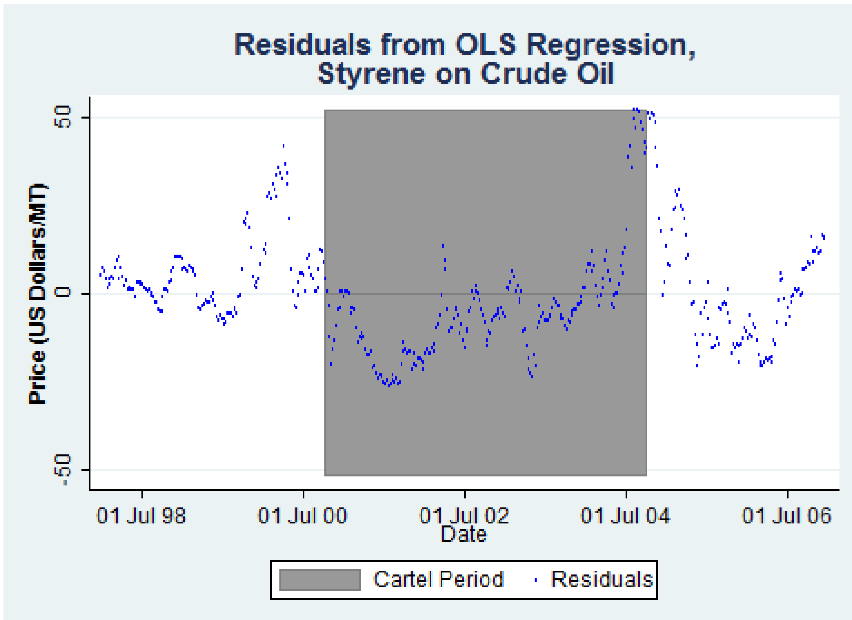

(a) Styrene-Crude OLS Residuals

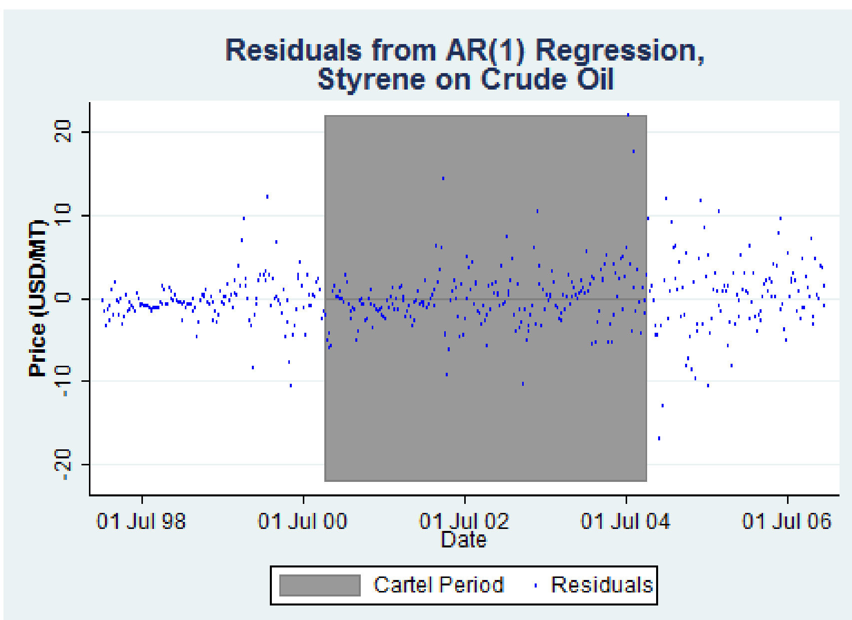

(b) Styrene-Crude AR(1) Residuals

Figure 3.2: Styrene-Crude Regression Residuals

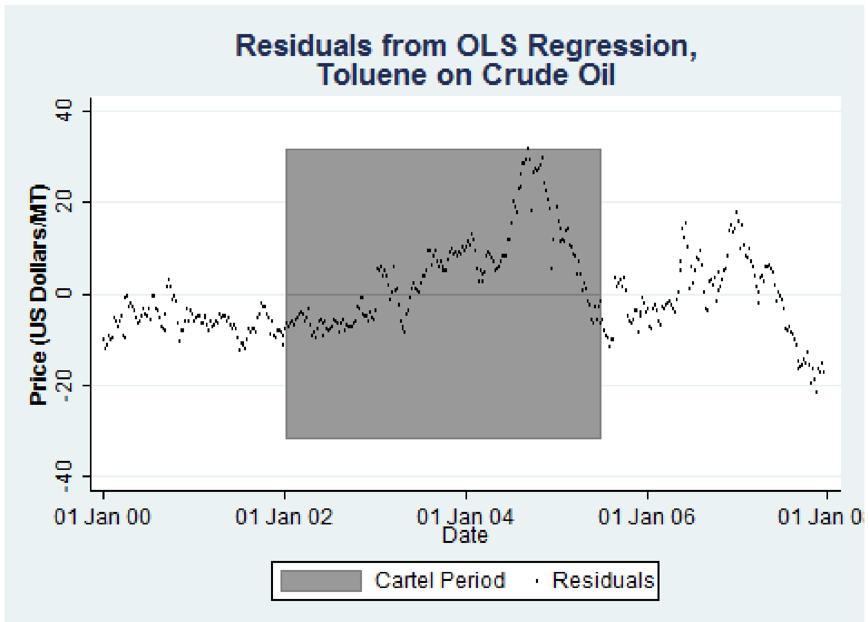

(a) Toluene-Crude OLS Residuals

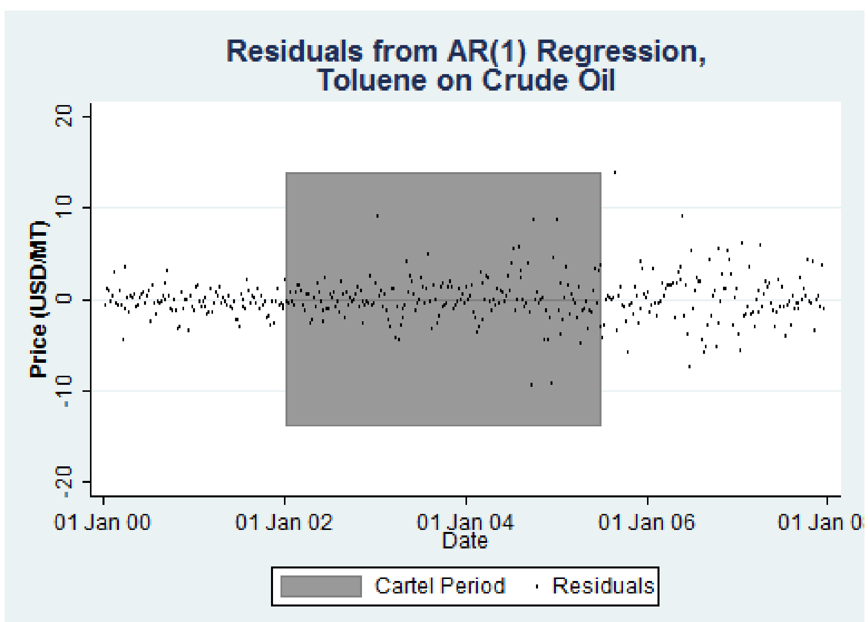

(b) Toluene-Crude AR(1) Residuals

Figure 3.3: Toluene-Crude Regression Residuals

the end of the sample. The largest residuals occur between July and September of 2004, and the variance gradually decreases over the rest of the sample period. We therefore treat the hypothesized "break date" of the cartel as September of 2004 for both products.

3.4 Empirical Specification

This section builds on the descriptive statistics presented in section 3.3.2 to estimate parameters of price asymmetry in pre- and post-collusion periods. It specifies a standard model of error correction based on Bachmeier and Griffin (2003) amd Borenstein, Cameron and Gilbert (1997). The model is estimated in two stages, as detailed by Engle and Granger (1987).[21]

$$P_t = \alpha_0 + \alpha_1 C_t + \delta_q + \epsilon_t \tag{3.1}$$

$$\begin{aligned}\Delta P_t = &\sum_{j=0}^{n} \left(\beta^+ \Delta C^+_{t-k} + \beta^- \Delta C^-_{t-k}\right) + \sum_{k=1}^{n} \left(\beta^+ \Delta P^+_{t-k} + \beta^- \Delta P^-_{t-k}\right) \\ &+ \rho \left(P_{t-1} - \delta_q - \alpha_0 - \alpha_1 C_{t-1}\right) + \mu_t \end{aligned} \tag{3.2}$$

The first stage, eq. (3.1), models output (petrochemical) price as a function of contemporaneous crude oil price and quarterly dummy variables. Equation (3.2) regresses the first-differenced output price on first differences of crude oil and output price, as well as on the lagged residuals from eq. (3.1). The positive and negative superscripts represent lagged differences with positive and negative signs, respectively. Comparing the absolute value of the magnitude of complementary coefficients constitutes a test of asymmetric pricing. With perfectly symmetric prices, the absolute value of each positive and negative coefficient.

The model relies on two identification assumptions. First, styrene, toluene, and crude oil price are assumed to be non-stationary and and integrated of order n. Consistent with the autoregressive depictions in fig. 3.2 and fig. 3.3, $n = 1$. Second, the joint distribution of P_t and C_t is cointegrated of order one. This assumption is consistent with a Dickey-Fuller test for crude oil price that rejects the null hypothesis of no unit root, together with the first assumption.[22]

[21] Other estimation approaches for vector error-correction models include simultaneous estimation through nonlinear least squares (Lewis (2011)) and nonlinear Bayesian techniques (Verlinda (2008)).

[22] $p < 0.05$ with one lag; $p > 0.10$ with two lags.

Table 3.2: VEC(1) Specification, Pre- and Post-Cartel Break

	Styrene		Toluene	
	(1) Pre-Break	(2) Post-Break	(1) Pre-Break	(2) Post-Break
ΔC_t^+	0.235	-0.280	0.452***	0.338
	(0.273)	(0.411)	(0.162)	(0.223)
ΔC_t^-	-0.043	-1.388***	-0.088	-0.784***
	(0.207)	(0.433)	(0.114)	(0.239)
ΔC_{t-1}^+	-0.023	1.062**	0.163	0.343
	(0.270)	(0.412)	(0.162)	(0.227)
ΔC_{t-1}^-	-0.675***	0.513	-0.346***	-0.097
	(0.211)	(0.450)	(0.116)	(0.246)
ΔP_{t-1}^+	0.188**	-0.058	0.075	-0.040
	(0.079)	(0.169)	(0.096)	(0.124)
ΔP_{t-1}^-	-0.480***	-0.384**	-0.475***	0.098
	(0.110)	(0.155)	(0.129)	(0.147)
ρ_{t-1}	-0.008	-0.038*	0.002	-0.069***
	(0.014)	(0.023)	(0.015)	(0.025)
N	345	116	241	168

***$p < 0.01$, **$p < 0.05$, *$p < 0.1$. The estimates are from the second stage of a two-step vector error correction model. The first stage includes quarter dummy variables; intercept of second stage not reported.
† Denotes confidence level from Wald Test between C_t^+ and $C_t^=$

Table 3.2 shows the results from the two-step error correction model, separately by chemical and by pre/post cartel period. Several coefficients are consistent with the specification of the model. The error correction term, ρ_{t-1}, is negative and significant. This is consistent with prices "drifting" back to the relationship estimated by $\hat{\alpha}_0$ and $\hat{\alpha}_1$ in eq. (3.1). Moreover, the estimated coefficients for cost increases and decreases have the expected signs in most periods.

There are two strong relationships to note among the remaining coefficient estimates. First, both contemporaneous cost coefficients display higher absolute value for positive differences than negative differences during the pre-break period. The toluene coefficients can be rejected for significance at the 1% level. The relationship flips in the post-break period, however: negative cost differences display higher absolute values than positive cost differences. In fact, column (2) in the styrene regression shows that the sign of the positive cost difference coefficient flips altogether: positive cost shocks were associated with negative styrene shocks, conditional on the lagged first differences, the error correction term, and the quarter dummys. The results reinforce the graphical depictions of an inversion of the "rockets and feathers" relationship to a "feathers and rockets" relationship, particularly for the styrene product.

3.5 Conclusion

The results outlined above have demonstrated a clear pattern that emerged in the Korean petrochemicals cartel under study. These relationships are evident in both descriptive statistics as well as more structured regression specifications. First, the degree of price-cost asymmetry between each product and crude oil increased during the cartel period relative to after. Second, the price-cost asymmetry inverted following the breakdown of the cartel, with adjustment ratios less than one in styrene and significantly different absolute values of coefficient magnitudes in styerne and toluene. In addition, the linear correlation between output and input price decreased following breakdown, and the time-series variation in output price was larger during the cartel period than after the cartel period.

The results are consistent with firms in the styrene market adhering to the collusive strategy disclosed during the KFTC's investigation. One avenue for further research is to examine the conditions under which this strategy constitutes optimal pricing in the presence of an active antitrust authority. In particular, the relationship between price and cost that arises during collusion, as well as the resulting breakdown in the relationship after the collapse of collusion, is consistent with the model outlined by Harrington and Chen (2006) and its precursors.[23] The descriptive figures imply that the price-cost margin increased gradually during the cartel period, but decreased suddenly upon cartel breakdown.

The evidence in this paper provides a basis for developing an empirical screen for collusion in fuel-intensive industries based on the degree of price-cost asymmetry. These measures were shown to provide more evidence consistent with a theory of collusive pricing than corresponding measures based only on the first and second moments of the picing distribution. The analysis relied on market-wide prices for petrochemicals as well as publicly available, international price series for crude oil. This data is readily available to competition authorities around the world and does not require firm-specific price or output information.

[23] See Harrington (2004) and Harrington (2005).

APPENDIX A

Appendix for Chapter I

A.1 Proofs of Results

A.1.1 Proof of Proposition 1

Proof 1.

Part I

Period two punishment from deviating in period one is weaker than period three punishment from deviating in period two.

$$q_{i1}^{*D1} > q_{i1}^{*} > q_{i1}^{*J1} \Rightarrow \Pi_{i2}^{*J1} - \Pi_{i2}^{*D1} = \underbrace{\frac{\left(a - c_1 + \lambda q_{i1}^{*J1}\right)^2}{8b}}_{< \frac{(a-\underline{c})^2}{8b}} - \underbrace{\frac{(a - c_1 + \lambda q_{i1}^{*D1})^2}{9b}}_{= \frac{(a-\underline{c})^2}{9b} \text{ if } c_1 - \lambda q_{i1}^{*D} = \underline{c}} < \frac{(a-\underline{c})^2}{72b}$$

If i deviates enough to reach $\underline{c}$ in period two, the inequality holds.

If i deviates but not enough to reach $\underline{c}$ in period two, the inequality still holds:

$$\frac{(a - c_1 + \lambda q_{i1}^{*D1})^2}{9b} - \frac{(a-\underline{c})^2}{9b} > 0 \equiv A$$
$$\frac{(a - c_1 + \lambda q_{i1}^{*J1})^2}{8b} - \frac{(a-\underline{c})^2}{8b} > 0 \equiv B$$
$$\frac{\partial A}{\partial \underline{c}} > \frac{\partial B}{\partial \underline{c}}$$

Part II

The same-period gain to deviating at period two is not boundlessly larger than the same-period gain to deviating at period one.

$$\frac{(a-\underline{c})^2}{72b} - \left(\frac{\left(a-c_1+\lambda q_{i1}^{*J1}\right)^2}{8b} - \frac{(a-c_1+\lambda q_{i1}^{*D1})^2}{9b}\right) >$$

$$\frac{1}{\delta}\left(\frac{(a-c_1)+\lambda q_{i1}^{*}}{64b} - (\Pi_{i1}^{*D1} - \Pi_{i1}^{*J1})\right)$$

$$\delta \underbrace{|\Delta_{1\to 2}\{\text{Deviation Punishment}\}_{t+1}|}_{\text{LHS}} > \underbrace{|\Delta_{1\to 2}\{\text{Deviation Gain}\}_{t}|}_{\text{RHS}}$$

$$\delta_1^{*J1} > \delta_2^{*J2} \quad \blacksquare$$

A.1.2 Proof of Proposition 2

Proof 2.

$$\delta_1^{*J1} < \delta_2^{*J2} < \delta \;\Rightarrow$$

$$\theta_1 = \theta_2 = 1 \tag{A.1}$$

$$\left[\frac{|\eta_1|}{|\eta_1|-1} - \frac{|\eta_1|}{|\eta_1|-\frac{1}{2}}\right] > \left[\frac{|\eta_2|}{|\eta_2|-1} - \frac{|\eta_2|}{|\eta_2|-\frac{1}{2}}\right] \qquad \forall\,|\eta_t|>1 \tag{A.2}$$

$$\frac{\partial C_i^*}{\partial q_{i1}}\left[\frac{|\eta_1|}{|\eta_1|-1} - \frac{|\eta_1|}{|\eta_1|-\frac{1}{2}}\right] > \frac{\partial C_i}{\partial q_{i2}}\left[\frac{|\eta_2|}{|\eta_2|-1} - \frac{|\eta_2|}{|\eta_2|-\frac{1}{2}}\right] \qquad \forall\,|\eta_t|>1 \tag{A.3}$$

$$\frac{\partial C_i^*}{\partial q_{i1}}\left[\frac{|\eta_1|}{|\eta_1|-1}-\frac{|\eta_1|}{|\eta_1|-\frac{1}{2}}\right]+\underbrace{\left(\frac{\partial C_i^{*J1}}{\partial q_{i1}}-\frac{\partial C_i^*}{\partial q_{i1}}\right)\frac{|\eta_1|}{|\eta_1|-1}}_{>0}>$$
$$\frac{\partial C_i}{\partial q_{i2}}\left[\frac{|\eta_2|}{|\eta_2|-1}-\frac{|\eta_2|}{|\eta_2|-\frac{1}{2}}\right]\forall\,|\eta_t|>1 \tag{A.4}$$

$$\Delta P_1 > \Delta P_2 \tag{A.5}$$

Inequality A.2 follows from the assumption that $|\eta_1| \geq |\eta_2|$: demand is at least as elastic in period two as in one. Inequality A.3 follows from inequality A.2 and the fact that $\frac{\partial C_i}{\partial q_{i1}} > \frac{\partial C_i}{\partial q_{i2}}$. Dynamic marginal cost is higher at period one, when static marginal costs are higher and the learning effect is present, than at period two, when static marginal costs are lower and the learning effect is not present.[1]

$$\delta_1^{*J1} < \delta < \delta_2^{*J2} \Rightarrow$$
$$\frac{1}{2} \leq \theta_1 < \theta_2 = 1 \tag{A.6}$$

$$\frac{\partial C_i^*(\theta_1)}{\partial q_{i1}}\left[\frac{|\eta_1|}{|\eta_1|-\theta_1}-\frac{|\eta_1|}{|\eta_1|-\frac{1}{2}}\right]<\frac{\partial C_i}{\partial q_{i2}}\left[\frac{|\eta_2|}{|\eta_2|-1}-\frac{|\eta_2|}{|\eta_2|-\frac{1}{2}}\right]$$
$$\forall\,|\eta_t|>1,\ |\eta_2|-|\eta_1|<\epsilon,\ \frac{\partial C_i^*(\theta_1)}{\partial q_{i1}}-\frac{\partial C_i}{\partial q_{i1}}<\gamma \tag{A.7}$$

$$\Delta P_1 < \Delta P_2 \qquad \forall\,|\eta_t|>1,\ |\eta_2|-|\eta_1|<\epsilon,\ \frac{\partial C_i^*(\theta_1)}{\partial q_{i1}}-\frac{\partial C_i}{\partial q_{i1}}<\gamma \tag{A.8}$$

$$\Delta P_1 > \Delta P_2 \qquad \forall\,|\eta_t|>1,\ |\eta_2|-|\eta_1|>\epsilon\text{ , or }\frac{\partial C_i^*(\theta_1)}{\partial q_{i1}}-\frac{\partial C_i}{\partial q_{i1}}>\gamma \quad \blacksquare \tag{A.9}$$

The intuition is similar in the above scenario, where conditions for $\Delta P_1 < \Delta P_2$ are

[1] At the first period, $\frac{\partial C_i^{*J1}}{\partial q_{i1}} \geq \frac{\partial C_i^*}{\partial q_{i1}}$: dynamic marginal cost is higher for collusion than competition. This occurs because firms in collusion plan to cut period two output relative to the competitive rate, and lower period two output increases the effective cost of period one output by diminishing its learning advantage.

derived. The key step is A.6: if the discount factor is high enough to sustain collusion at the joint profit maximizing level at period two, but not at period one, then $\theta_2 = 1$ and $\theta_1 < 1$. The remainder of the proof shows that if this difference in the average strength of collusion is high enough relative to the evolution of dynamic marginal cost and of price elasticity, then $\Delta P_1 < \Delta P_2$. If not, then $\Delta P_1 > \Delta P_2$.

A.2 Learning Spillovers in Collusion

The model presented in section 1.2 treats learning entirely as a proprietary possibility: firms reduce future cost as a function of their own current output, but not as a function of others' output. Learning diffusion between firms is not tied to any firm's output. Theory and empirical applications in learning-by-doing industries raise the possibility that one firm could also benefit from another firm's output—learning "spills over" between firms.[2] This could happen if learning is human capital-specific and knowledgable employees change firms, or if research is published in scientific journals once findings are made.

This section extends the baseline model by allowing learning to spill over exogenously between firms. I show that the mechanism that makes collusion less sustainable during the learning period—one firm gains a strategic advantage over the other if it deviates from an output restriction during the first period—continues to hold under plausible spillover rates. Specifically, as long as the rate at which a firm reduces its own cost is greater than the rate at which it reduces its competitor's cost, firms play as strategic substitutes. The sufficient conditions for Results 1.1 and 1.2 then extend to the learning spillovers case as well: collusion requires more patience to sustain when firms begin colluding during the learning phase relative to the post-learning or diffusion phase.

A.2.1 Spillover Effect and Strategic Substitutability

Consider the model presented in section 1.2.1: firms i and j compete in an infinite horizon game with three distinct stages. In period one, they have identical marginal costs; by period two, costs decline based on period one's output; and by period τ, which is infinitely repeated, costs reach their minimum level and there is no further learning. The possibility of spillover-free diffusion remains: if firms do not learn enough to exhaust learning by the

[2] Fudenberg and Tirole (1983) and Ghemawat and Spence (1985) discuss spillovers in their learning-by-doing models. Zulehner (2003) and Irwin and Klenow (1994) use structural models to infer the existence of learning spillovers in DRAM based on firms' output reactions.

start of period two, further cost reductions diffuse freely between firms.

Introduce the parameter $0 < \alpha < 1$ to represent the rate of cost reduction in c_{j2} based on output q_{i1}.[3] Firm i's baseline first order condition in the competitive case becomes:

$$\begin{aligned}
\underset{q_i}{max}\ \Pi_i &= [a - b(q_{i1} + q_{j1}) - c_1]q_{i1} + \delta[a - b(q_{i2} + q_{j2}) - (c_1 - \lambda q_{i1} - \alpha q_{j1})]q_{i2} \\
&\quad + \sum_{\tau=3}^{\infty} \delta^{\tau-1}[a - b(q_{i\tau} + q_{j\tau}) - \underline{c}]q_{i\tau} \\
&\Rightarrow \frac{\partial \Pi_i}{\partial q_{i1}} :\ a - 2bq_{i1} - bq_{j1} - c_1 + \underbrace{(\delta\lambda q_{i2})}_{\text{“PLE”}} - \underbrace{\left(\delta b q_{i2} \frac{\partial q_{j2}}{\partial q_{i1}}\right)}_{\text{“ISE”}} = 0 \qquad \text{(A.10)}
\end{aligned}$$

The first order condition is identical to the spillover-free case, which includes the proprietary learning incentive (“PLE”), except for one additional term. The intertemporal strategic parameter (“ISE”) captures the total change in firm j's output in period two that is attributable to a change in firm i's output in period one. Its effect depends on the sign of $\frac{\partial q_{j2}}{\partial q_{i1}}$: if $\frac{\partial q_{j2}}{\partial q_{i1}} > 0$, then the ISE runs counter to the PLE. Intuitively, this happens when spillovers are relatively large: i reduces j's period two costs via the spillover parameter α. The PLE induces firms to produce more than static marginal cost through the direct effect of learning; the ISE can induce firms to cut back their extra output through the indirect effect of learning on their rivals.[4] The net effect of the PLE and ISE determines whether firms play as strategic substitutes, as in the learning-free case, or strategic complements, when spillovers prevent firms from producing the “extra” output that learning encourages.

[3]Like proprietary learning parameter λ, spillover parameter α is identical between firms i and j. The game therefore remains symmetric.

[4] $\frac{\partial q_{j2}}{\partial q_{i1}}$ is not positive for all values of α because the ISE also takes into account the inherent strategic substitutability of Cournot competiton. If $\alpha = 0$, for example, then the ISE induces firms to produce even more in the first period than they would in its absence.

A.2.2 Spillover Effect and Incentive Compatibility

To assess the role of learning spillovers in collusion, standardize and fix total learning to one: ($\lambda + \alpha = 1$). The ability to restrict output during coooperative play depends on i's incentive to deviate. i's incentive to deviate at period one depends on the profits it anticipates receiving in periods two.[5] Second period profits depend on second period outputs, which in turn depend on the value of the ISE. Because i considers its deviation decision holding j's output fixed at the cartel level, second period outputs can be expressed through the first order conditions as:

$$q_{i2} = \frac{a - bq_{j2} - c_1 + \lambda q_{i1} + \alpha q_{j1}^{*J1}}{2b} \tag{A.11}$$

$$q_{j2} = \frac{a - bq_{i2} - c_1 + \lambda q_{j1}^{*J1} + \alpha q_{i1}}{2b} \tag{A.12}$$

Substituting A.11 into A.12 yields:

$$q_{j2} = \frac{a - c_1}{3b} + \frac{q_{i1}(2\alpha - \lambda)}{3b} + \frac{q_{j1}^{*J1}(2\lambda - \alpha)}{12b} \tag{A.13}$$

$$\Rightarrow \frac{\partial q_{j2}}{\partial q_{i1}} = \frac{2\alpha - \lambda}{3b} \tag{A.14}$$

Equation A.14 illustrates the role of spillovers in determining the sign of the ISE. $\frac{\partial q_{j2}}{\partial q_{i1}} > 0$ when $\alpha > \frac{1}{3}$: spillovers discourage i from deviating when they reduce cost at least half as much as does proprietary learning. Whether firms play as strategic substitutes or complements, however, depends on the net effect of the ISE and PLE shown in i's first order condition above. Substituting A.14 into A.10 reveals the central relationship governing

[5] It also depends on τ^{th} period profits. As shown in the ICC's **??** and **??**, τ^{th} period profits are identical whether collusion begins at period one or at period two.

strategic competition:

$$a - 2bq_{i1} - bq_{j1} - c_1 + \underbrace{\delta q_{i2} \left(\frac{2(\lambda - \alpha)}{3} \right)}_{\text{PLE + ISE}} = 0 \tag{A.15}$$

$$i \text{ and } j \text{ play as: } \begin{cases} \text{strategic substitutes} & \text{if } \lambda > \frac{1}{2} \\ \text{strategic complements} & \text{if } \lambda < \frac{1}{2} \end{cases}$$

The inequality above states that if the proprietary learning rate is greater than the spillover learning rate, i and j will compete as strategic substitutes, just as in the learning-free case. Strategic substitutability ensures that the punishment for deviating at period one is weaker than the punishment for deviating at period two. The sufficient condition for Result 1.1 therefore applies to the learning spillovers case if $\alpha < \lambda$. Whenever the dynamic effect of strategic substitution dominates any static effect of decreased elasticity, the joint profit maximizing collusive equilibrium is sustainable at a lower minimum discount factor starting from period two (the post-learning or diffusion period) than from period one (the learning period). The subsequent implications for price differences between generations outlined in Result 1.2 also follow.

A.2.3 Discussion of Empirical Spillover Rates

In practice, it is unlikely that learning spillovers are large enough to render the quantity-setting game one of strategic complements, even when the spillover effect is magnified by $n > 2$ firms in the market. Zulehner (2003) finds that if spillovers are large enough to induce strategic complementarity, complementarity occurs only at the end of the product life cycle, when large learning effects have subsided. Such a result is fully consistent with this paper's findings, because increased spillovers at the end of the product life cycle would reinforce collusion for older generations but not for newer generations. The conclusions that follow from the plant-level DRAM manufacturing analyses of Hatch and Mowery (1998) and Macher and Mowery (2003), moreover, imply that it is more likely that the learning spillover effects found in empirical DRAM papers are conflated with the effect of general cost diffusion over the product life cycle. Cost diffusion is modeled through the transition between periods two and τ in the framework presented in Section 1.2.

A.3 DRAM Output by Technology

Section 1.3 describes the rival technologies available to DRAM manufacturers during the early 2000s. DDR, developed cooperatively by DRAM fabricators through JEDEC, faced a rival technology in RDRAM, licensed by the design firm Rambus. Although RDRAM boasted slightly faster performance speeds than DDR, the two technologies were largely substitutable. After initial backing from Intel, RDRAM failed to gain popularity among fabricators and OEM's and DDR quickly became the industry standard upon which subsequent innovations have been based.

The empirical pattern of positive collusive overcharge for the 16Mb and 64Mb generations but zero or negative overcharges for the 128Mb of 256Mb generations cannot be explained by changes in the underlying technology of the chips. During the 1990s, the four largest firms—Samsung, Infineon, Micron and Hynix—designed their plants to accomodate both technology types for 128Mb and 256Mb chips, skipping the older two generations.[6] After DDR became the dominant next-generation technology, Rambus filed a private antitrust suit in 2004 accusing the DRAM manufacterers of flooding the market with DDR to price out RDRAM.[7]. Note that this case is distinct from both the DOJ antitrust case around which this paper focuses and the string of largely successful Rambus patent infringement suits that spurred the FTC and EC to file retaliatory "patent ambush" litigation. Figures A.1 and A.2 rule out the possibility that competition between firms pushing either of these technologies contributed to the output increase discussed in section 1.5.1, because RDRAM and DDR did not proliferate until after 2000:

[6]"Samsung's Fuller Disclosure," Businessweek 11/29/2005 http://www.businessweek.com/stories/2005-11-29/samsungs-fuller-disclosure

[7]The case proceeded to a $3.95 billion trial in 2011, where a jury ultimately rejected Rambus' allegation. See "Rambus Loses Antitrust Lawsuit, Shares Plunge," http://www.reuters.com/article/2011/11/16/us-rambus-micron-verdict-idUSTRE7AF1XL20111116

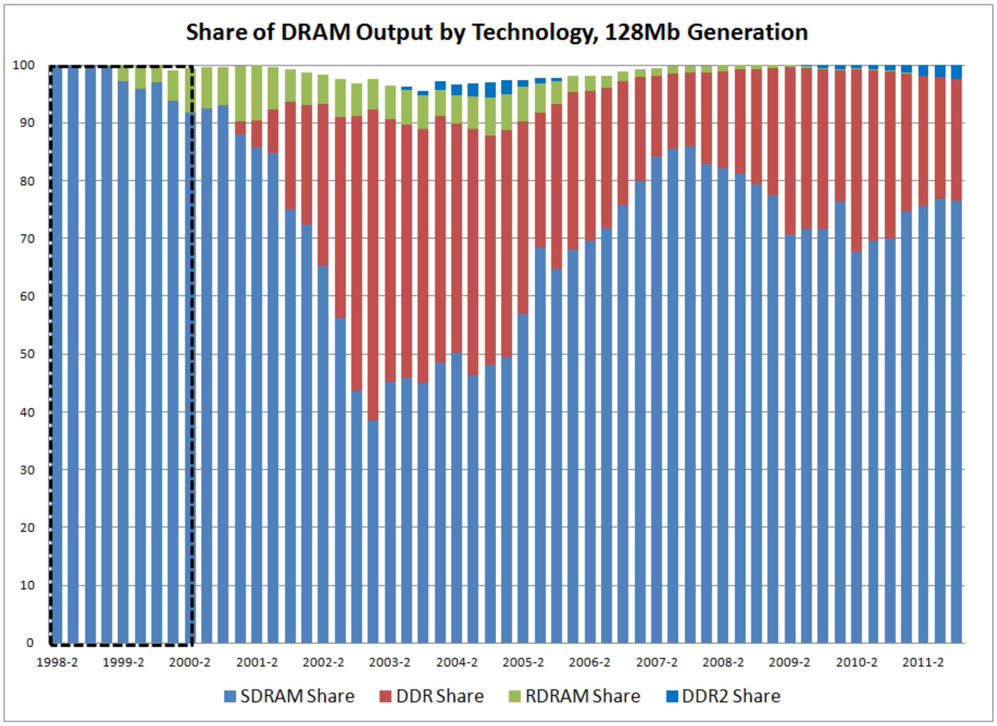

Figure A.1: 128Mb DRAM Technology Share, 1998-2011

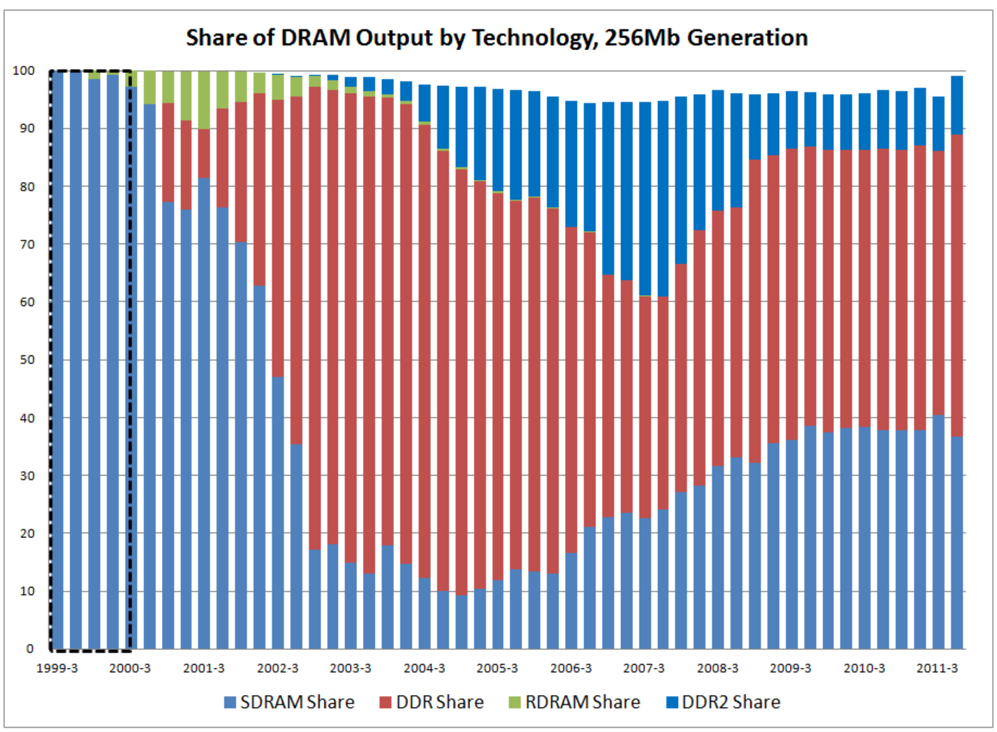

Figure A.2: 256Mb DRAM Technology Share, 1999-2011

A.4 Additional Tables

Table A.1: DRAM Years Active by Generation, 1974-2011

Generation	**Years Active***
4Kb[1]	1974-1985
16Kb	1976-1985
64Kb	1979-1995
256Kb	1982-1997
1Mb	1985-2002
4Mb	1988-2010
16Mb	1991-2011
64Mb	1996-2011
128Mb	1998-2011
256Mb	1999-2011
512Mb	2001-2011
1Gb	2003-2011
2Gb	2005-2011
4Gb	2010-2011

* Year is counted if at least one firm produces output in any quarter
[1] First year of Gartner data; shipments began about two years earlier

Table A.2: C_4 Concentration Ratios by Generation, 1994-2004*

Year	16Mb	64Mb	128Mb	256Mb
1994	60.0			
1995	51.4	100		
1996	51.4	76.3		
1997	42.8	79.8		
1998	53.4	56.2	100	
1999	62.6	67.6	75.1	100
2000	66.8	68.7	76.9	82.6
2001	68.5	73.6	74.3	74.2
2002	60.5	72.2	78.2	77.6
2003	73.7	69.2	79.2	78.4
2004	79.3	70.3	77.0	74.7

* The four largest firms across generations from 1994-2004 were Samsung, Infineon, Micron and Hynix.

In addition to changes in the market shares of the four pre-existing firms, year-to-year changes reflect new firms entering a generation and merger and acquisition activity. As characterized in Gartner Research's data: in 1998, Micron purchased Texas Instruments' DRAM operations. In 1999, SK Hynix (then named Hyundai) merged with LG Semiconductor. In 2001, NEC and Hitachi begain operating the DRAM joint venture Elpida Memory.

A.5 Additional Figures

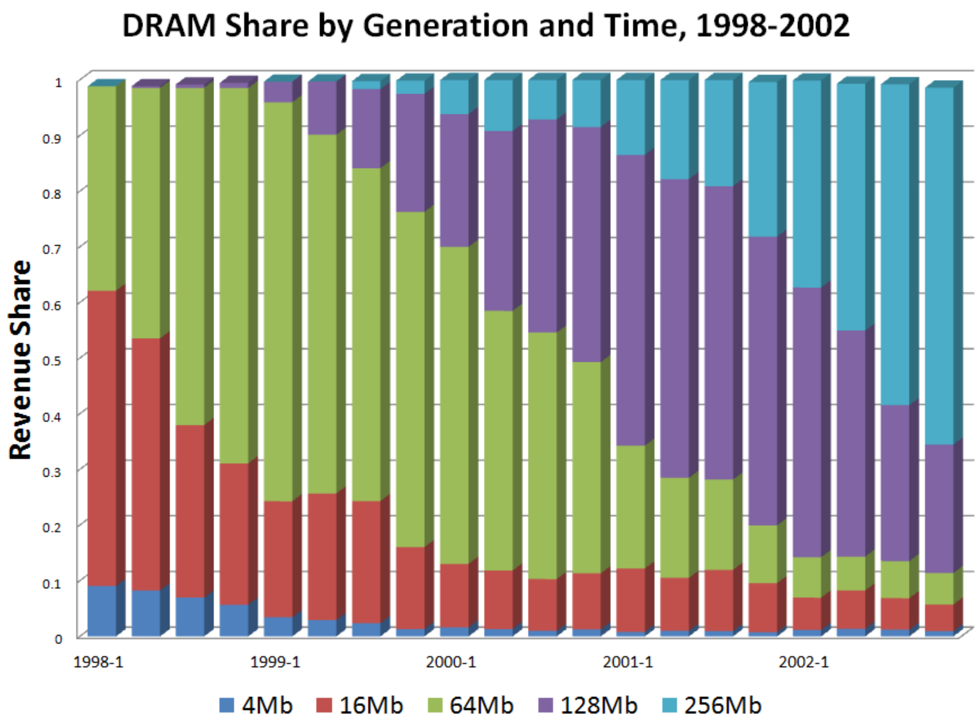

Figure A.3: Revenue Share: 4Mb - 256Mb

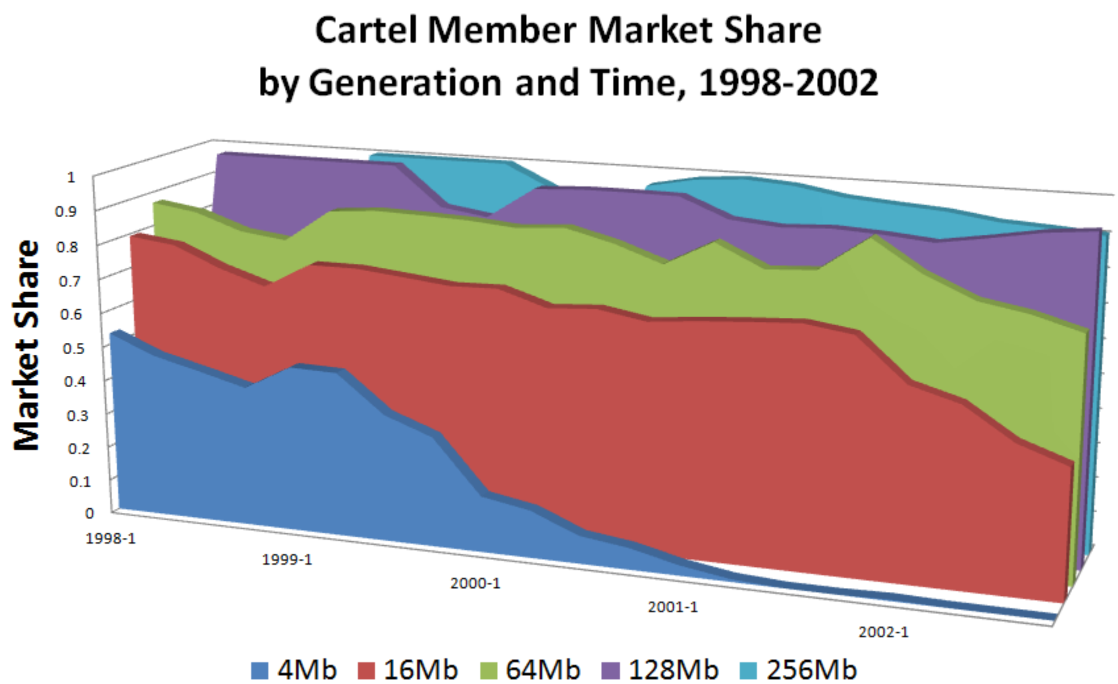

Figure A.4: Cartel Insider Market Share: 4Mb - 256Mb

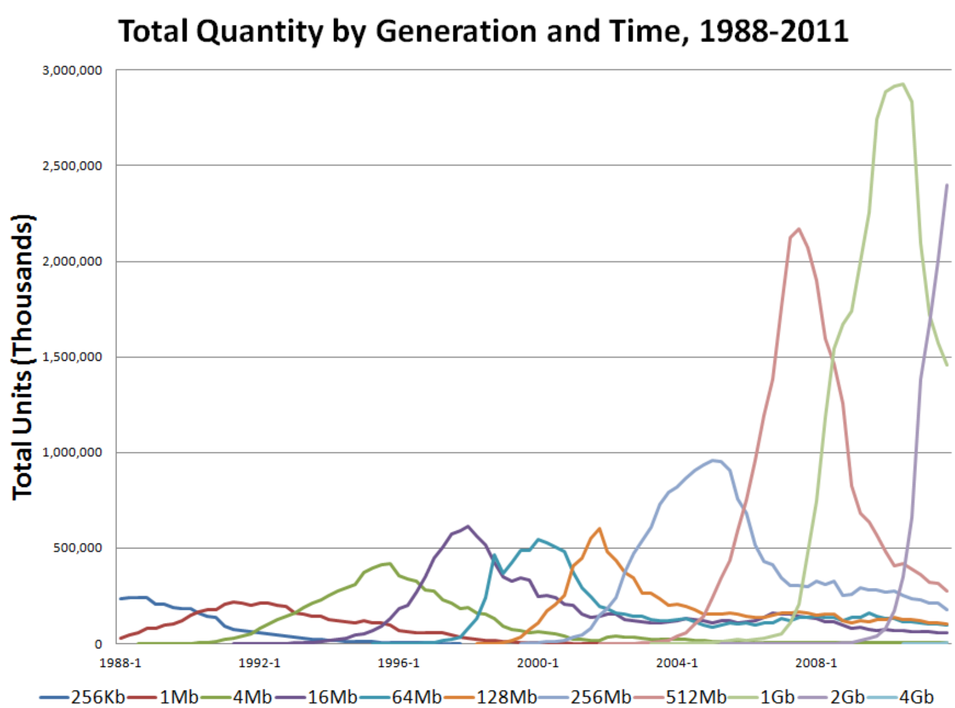

Figure A.5: Total Output by Generation: 256Kb - 4Gb

APPENDIX B

Appendix for Chapter II

B.1 Proofs of Results

B.1.1 Proof of Proposition 1

The sufficient and necessary condition guaranteeing that $\frac{Z_t^{Bs}}{Z_t^{B1}} > \frac{Z_t^{As}}{Z_t^{A1}}$ is stated as:

$$(J_2 - J_1)\left[\pi^P - \pi^D\right] > r_s(J_2) - r_s(J_1) \tag{B.1}$$

$$J_2\left[\pi^P - \pi^D\right] - r_s(J_2) > J_1\left[\pi^P - \pi^D\right] - r_s(J_1) \tag{B.2}$$

$$Pr\left[Fail_s\right]_t^B > Pr\left[Fail_s\right]_t^A \tag{B.3}$$

Equation (B.3) follows from condition 2.3: $Pr\left[Fail_t\right]^l = J\left[\pi^P - \pi^D\right] > r_s(J)$. The following inequalities relate the probability of a claim in staet l failing at time t to Z_t^{ls}, the number of claims from stage $s \in \{1, 2\}$ in state l that close in period t.

$$1 - Pr\left[Fail_s\right]_t^B < 1 - Pr\left[Fail_s\right]_t^A \tag{B.4}$$

$$X_t^{B1} < X_t^{A1} \tag{B.5}$$

$$Z_t^{B1} < Z_t^{A1} \tag{B.6}$$

$$\frac{Z_t^{Bs}}{Z_t^{B1}} > \frac{Z_t^{As}}{Z_t^{A1}} \quad \blacksquare \tag{B.7}$$

Where inequality B.6 is equivalent to inequality B.5 because the number of claims from an initial set of claims arriving at t that settle in t is equal to the set of claims that settle in t that are in stage one.

B.1.2 Proof of Proposition 2

Part i

The probability of failure for claim x at state $s \in \{1, 2\}$ in period t can be characterized as:

$$\pi^P - \pi^D > \frac{r_s(J)}{J} \tag{B.8}$$

$$\pi^P \left(Y', \epsilon(\sigma_t^{2P})\right) - \pi^D \left(Y', \epsilon(\sigma_t^{2D})\right) > \frac{r_s(J)}{J} \tag{B.9}$$

Where $\hat{Y}^i = Y' + \epsilon,\ \epsilon \sim N(0, \sigma^2)$: subjective probabilities are formed as a function of the normally distributed error term ϵ.

This probability is composed of two mutually exclusive and nonstochastic outcomes: the parties fail to settle at stage s, or they settle at state s. Let K^* denote a constant. Then settlement failure and settlement imply two inequalities:

$$\sigma_t^{2P} - \sigma_t^{2D} > K^* > 0 \Rightarrow\ x \text{ fails to settle at } s \tag{B.10}$$

$$\sigma_t^{2P} - \sigma_t^{2D} < 0 < K^* \Rightarrow x \text{ settles at } s \tag{B.11}$$

The sufficient condition of Proposition 2.2 (i) states that $\sigma_t^2 > \sigma_{t+1}^2$, with the mean centered at zero in both t and $t + 1$. Letting $\alpha > 0$, this implies:

$$E[\sigma_{t+1}^{2P} - \sigma_{t+1}^{2D} | \sigma_{t+1}^{2P} - \sigma_{t+1}^{2D} > 0] < E[\sigma_t^{2P} - \sigma_t^{2D} | \sigma_t^{2P} - \sigma_t^{2D} > 0] \tag{B.12}$$

$$E\left[\sigma_{t+1}^{2P} - \sigma_{t+1}^{2D} | \sigma_t^{2P} - \sigma_t^{2D} > 0\right] = \alpha E\left[\left(\sigma_t^{2P} - \sigma_t^{2D} | \sigma_{t+1}^{2P} - \sigma_{t+1}^{2D} > 0\right)\right] \tag{B.13}$$

Mapping the discrete events B.12 and B.13 back into stochastic probabilities,

$$E\left[\pi^P - \pi^D | \pi^P - \pi^D > 0\right]_t > E\left[\pi^P - \pi^D | \pi^P - \pi^D > 0\right]_{t+1} \tag{B.14}$$

$$Pr\left[Fail_s | \pi^P - \pi^D > 0\right]_t > Pr\left[Fail_s | \pi^P - \pi^D > 0\right]_{t+1} \tag{B.15}$$

When parties obtain mutually optimistic draws from ϵ, then the average magnitude of disagreement rises and the probability of disagreement rises with it.

Observe, however, that the corresponding inference cannot be made for the discrete condition that $\sigma^{2P} - \sigma^{2D} > 0$. This is because when parties obtain mutually pessimistic draws from ϵ, the probability that they fail to settle is zero: $\sigma_t^{2P} - \sigma_t^{2D} < 0 \Rightarrow \sigma_t^{2P} - \sigma_t^{2D} < K^*$. This disparity is driven by strictly positive litigation costs $r_s(J)$: in the absence of optimism bias, parties possess no reason to reach disagreement.[1] Therefore

$$Pr\left[Fail_s | \pi^P - \pi^D < 0\right]_t = Pr\left[Fail_s | \pi^P - \pi^D > 0\right]_{t+1} = 0 \tag{B.16}$$

Finally, combining inequalities B.15 and B.16 implies the necessary condition:

$$X_t^{B1} > X_{t+1}^{A1} \tag{B.17}$$

$$Z_t^{B1} > Z_{t+1}^{A1} \tag{B.18}$$

$$\frac{Z_t^{Bs}}{Z_t^{B1}} < \frac{Z_{t+1}^{Bs}}{Z_t^{B1}} \quad \blacksquare \tag{B.19}$$

Part ii

The proof of (ii) is similar to the proof of Proposition 2.1 demonstrated in appendix B.1.1 above. Let $0 < r_s(J_2) - r_s(J_1) < (J_2 - J_1)\left(\pi^P - \pi^D\right)$ in period $t + 1$. Because

[1] This standard result of Divergent Expectations models is *ceteris paribus*. In particular, it does not hold when parties have asymmetric stakes in the trial.

$\sigma_{t+1}^{2B} = \sigma_{t+1}^{2A}$, the following implication holds:

$$\left[F^P\left(Y'+\epsilon\right)-F^D\left(Y'+\epsilon\right)\right]_{t+1}^{B}-\frac{r_s\left(J\right)}{J}>\left[F^P\left(Y'+\epsilon\right)-F^D\left(Y'+\epsilon\right)\right]_{t+1}^{A}-\frac{r_s\left(J\right)}{J} \tag{B.20}$$

$$Pr\left[Fail_s\right]_{t+1}^{B}>Pr\left[Fail_s\right]_{t+1}^{A} \tag{B.21}$$

The step from B.20 to B.21 is proven by B.3 in Proposition 2.1.

Now let $J^{**}>\left(J_2-J_1\right)\left(\pi^P-\pi^D\right)$. Then

$$\left(J_2-J_1\right)\left(\pi^P-\pi^D\right)<J^{**}<r_s(J_2)-r_s(J_1) \tag{B.22}$$

$$Pr\left[Fail_s\right]_{t+1}^{B}<Pr\left[Fail_s\right]_{t+1}^{A} \tag{B.23}$$

$$\frac{Z_{t+1}^{Bs}}{Z_{t+1}^{B1}}>\frac{Z_{t+1}^{As}}{Z_{t+1}^{A1}}\ \blacksquare \tag{B.24}$$

Where the final step is the same as the one from B.3 to B.7.

B.2 Additional Tables

Table B.1: State Tort Liability Regimes
for Insurance Bad Faith to 1997[1,2,3]

State	Tort Start-End	State	Tort Start-End
AK	1974-	MO	
AL	1981-	MS	1984-
AR	1984-	MT	1982-
AZ	1982-	NC	1976-
CA	1973-	ND	1979-
CO	1983-	NE	
CT	1973-	NH	
DC		NJ	
DE	1982-1995	NM	1974-
FL	1975-1986	NV	1975-
GA	Unknown-1989	NY	
HI	1996-	OH	1983-
IA	1988-	OK	1977-
ID	1986-	OR	
IL		PA†	1990-
IN	1993-	RI	1980-
KS		SC	1983-
KY	1977-	SD	1986-
LA†	Unknown-	TN	
MA†	Unknown-	TX	1987-
MD		UT	
ME		VA	
MI		VT	1979-
MN		WA	1992-
		WI	1978-
		WV	
		WY	1990-

Based on Stempel (2006), Ostrager and Newman (2008), and Corp. and LLP (2008)
[1] Blank space indicates no punitive damages through 1997.
[2] "-" indicates punitive damages regime lasted through at least 1997. Illinois (2002) is only known state to have enacted punitive damages after 1997.
[3] "Unknown" is treated as pre-1977. Regression results robust to dropping these states.
† Punitive damages but no tort regime.

B.3 Additional Figures

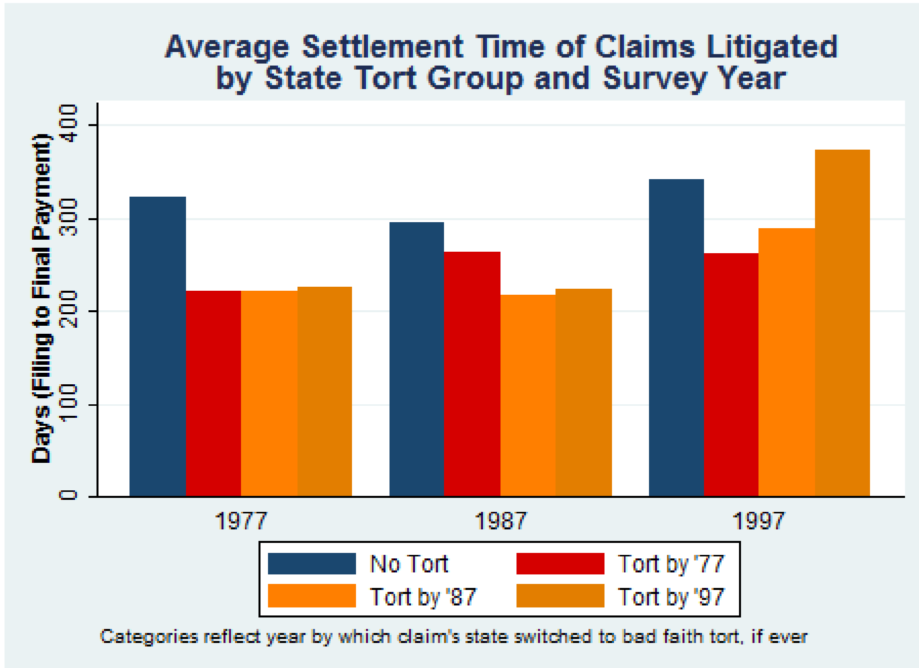

Figure B.1: Average Settlement Times by Tort Switch Group and Survey Year

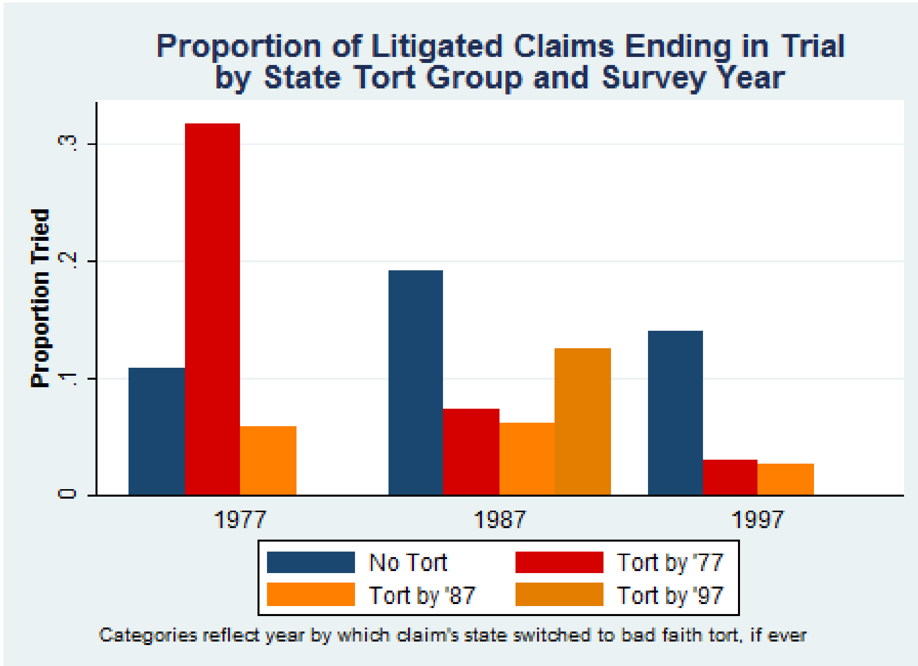

Figure B.2: Average Conditional Trial Rates by Tort Switch Group and Survey Year

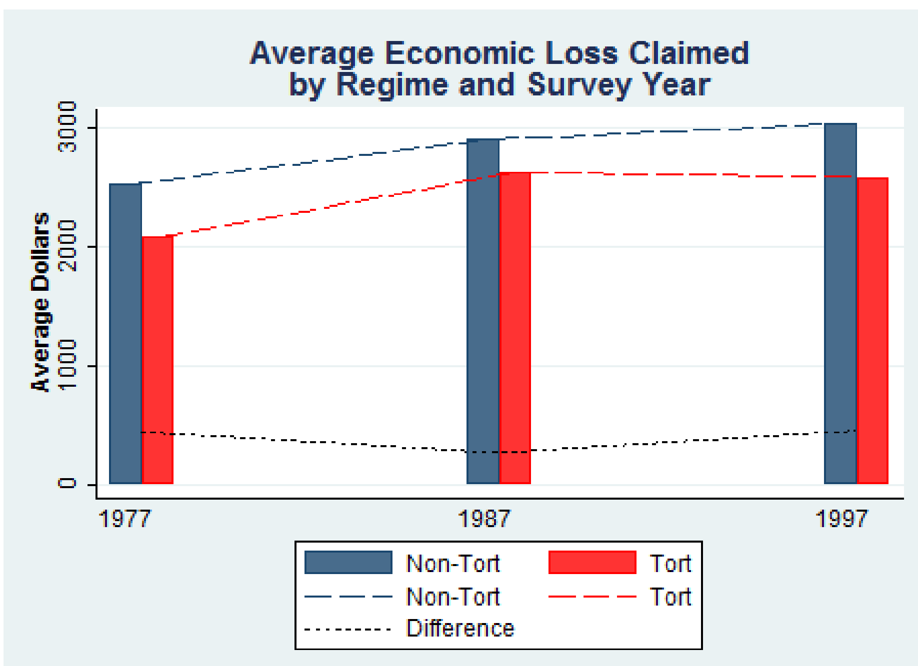

Figure B.3: Mean Economic Loss Claimed by Survey Year and Regime

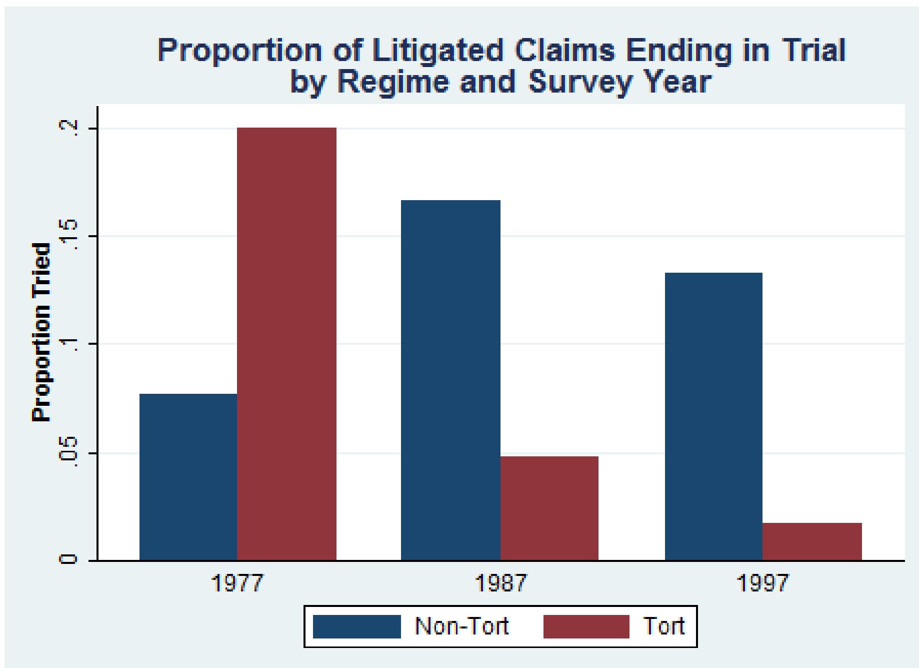

Figure B.4: Average Conditional Trial Rates, Excluding California

APPENDIX C

Appendix for Chapter III

C.1 Additional Tables

Table C.1: Korean Petrochemical Production
by Company Year 2007 (Thousands of Metric Tons)

Firm	Styrene	Toluene
Samsung Total†	870	
LG Chem	680	100
Honam Petrochem†	500	78
SK Energy†	384	849
BASF Company	320	
Yeochun NCC	290	233
Dongbu HiTek†	270	
S-Oil Corp.		350
GS Caltex†		170
OCI Company		30
Total	*3,314*	*1,810*

Source: Korean Petrochemical Industry Association
† Publicly fined by KFTC for price fixing

Table C.2: Aggregate Korean Petrochemical Production
and Export Percentage, Selected Years (Thousands of Metric Tons)

Product	**Styrene**		**Toluene**	
	Production (Thou/MT)	*Export %*	*Production (Thou/MT)*	*Export %*
1995	1,356	25	952	65
2000	2,430	37	1,440	70
2004	2,608	35	1,834	45
2005	2,722	35	1,792	42
2006	2,964	38	1,779	37

Source: Korean Petrochemical Industry Association

C.2 Additional Figures

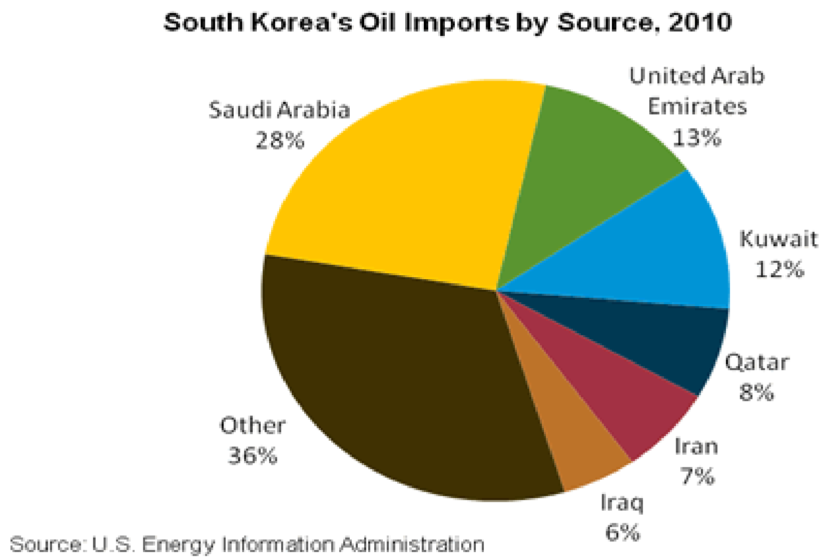

Figure C.1: Sources of Oil Import by Country, 2010

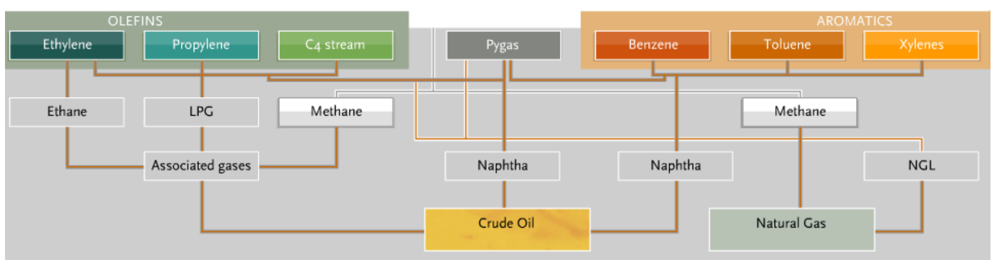

Figure C.2: Petrochemical Feedstock Diagram

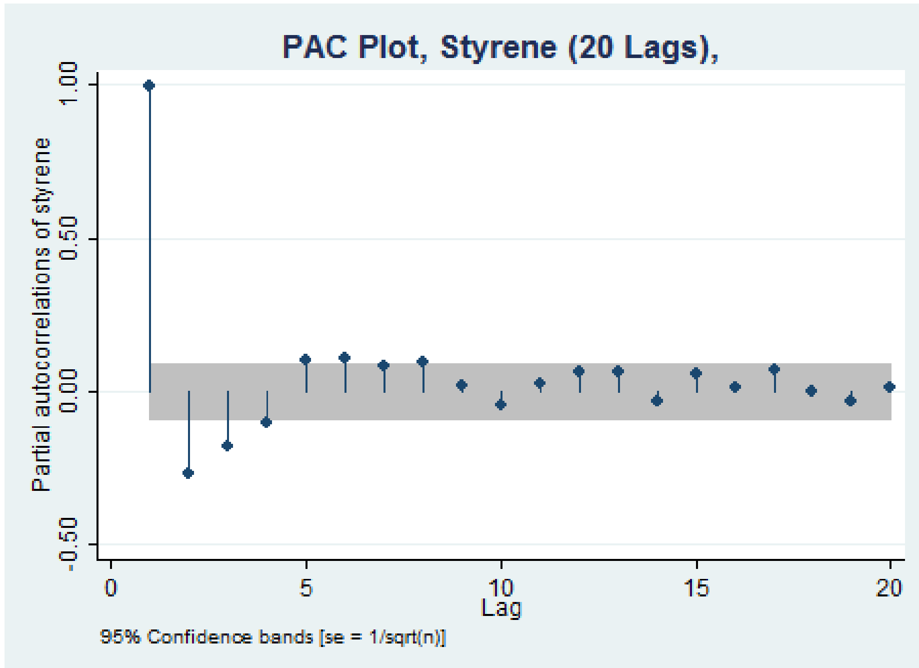

Figure C.3: Styrene Partial Autocorrelation Graph, 1998-2006

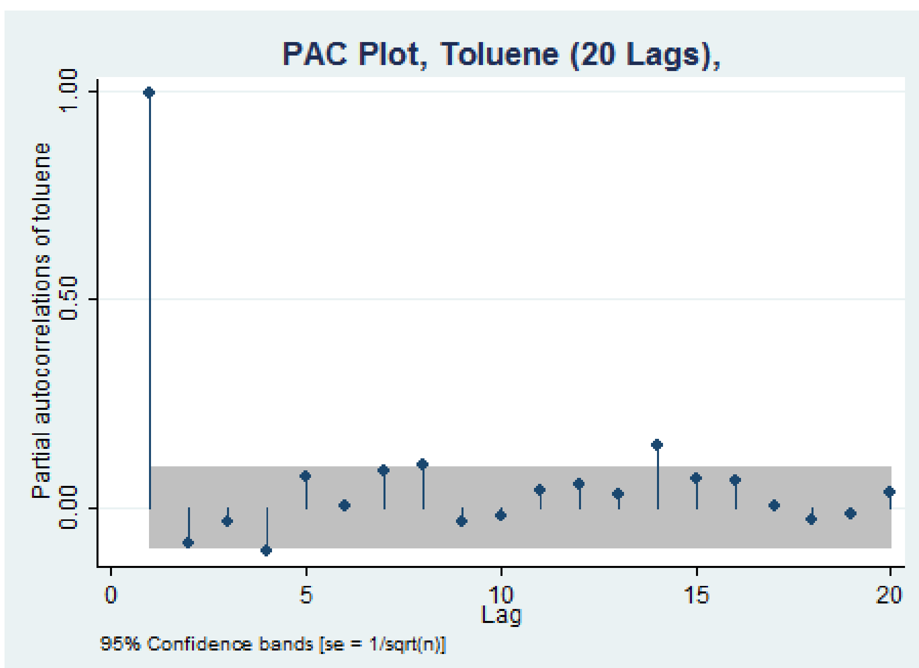

Figure C.4: Toluene Partial Autocorrelation Graph, 2000-2007

BIBLIOGRAPHY

Abraham, Kenneth S. 1994. "Natural History of the Insurer's Liability for Bad Faith." *Texas Law Review*, 72: 1295–1316.

Abrantes-Metz, Rosa M., Luke M. Froeb, John F. Geweke, and Christopher T. Taylor. 2006. "A Variance Screen for Collusion." *International Journal of Industrial Organization*, 24(3): 467–486.

Abreu, Dilip. 1986. "Extremal Equilibria of Oligopolistic Supergames." *Journal of Economic Theory*, 39(1): 191–225.

Abreu, Dilip. 1988. "On the Theory of Infinitely Repeated Games with Discounting." *Econometrica*, 56(2): 383–396.

Asker, John. 2010. "A Study of the Internal Organization of a Bidding Cartel." *American Economic Review*, 100(3): 724–762.

Asmat, Danial. 2015. "Collusion Along the Learning Curve: Theory and Evidence from the Semiconductor Industry." Working Paper, University of Michigan Ross School of Business.

Asmat, Danial, and Sharon Tennyson. 2014. "Does the Threat of Insurer Liability for "Bad Faith" Affect Insurance Settlements." *The Journal of Risk and Insurance*, 81(1): 1–26.

Asmat, Danial, and Sharon Tennyson. Forthcoming. "The Law and Economics of Insurance Bad Faith Liability." In *Research Handbook on the Economics of Insurance Law*. , ed. Daniel Schwarcz and Peter Siegelman. Elgar Publishing.

Bachmeier, Lance J., and James M. Griffin. 2003. "New Evidence on Asymmetric Gasoline Price Responses." *The Review of Economics and Statistics*, 85(3): 772–776.

Bajari, Patrick, and Lixin Ye. 2003. "Deciding Between Competition and Collusion." *The Review of Economics and Statistics*, 85(4): 971–989.

Bebchuk, Lucian A. 1984. "Litigation and Settlement Under Imperfect Information." *Bell Journal of Economics*, 15(3): 404–415.

Benoit, Jean-Pierre, and Vijay Krishna. 1987. "Dynamic Duopoly: Prices and Quantities." *Review of Economic Studies*, 54(1): 23–35.

Bernheim, Douglas, and Michael D. Whinston. 1990. "Multimarket Contact and Collusive Behavior." *Rand Journal of Economics*, 21(1): 1–26.

Besanko, David, Ulrich Doraszelski, Yaroslav Kryukov, and Mark Satterthwaite. 2010. "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics." *Econometrica*, 78(2): 453–508.

Bolotova, Yuliya, John M. Connor, and Douglas J. Miller. 2008. "The Impact of Collusion on Price Behavior: Empirical Results from two Recent Cases." *International Journal of Industrial Organization*, 26(6): 1290–1307.

Borenstein, Severin, Colin A. Cameron, and Richard Gilbert. 1997. "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" *Quarterly Journal of Economics*, 112(1): 305–339.

Bresnahan, Timothy F. 1989. "Empirical Studies of Industries with Market Power." In *The Handbook of Industrial Organization*. Vol. 2, , ed. Richard Schmalensee and Robert D. Willig, 1011–1057. Elsevier Science Publishers B. V.

Brod, Andrew, and Ram Shivakumar. 1999. "Advantageous Semi-Collusion." *The Journal of Industrial Economics*, 47(2): 221–230.

Browne, Mark J., and Joan T. Schmit. 2008. "Litigation Patterns in Automobile Bodily Injury Claims 1977-1997: Effects of Time and Tort Reforms." *The Journal of Risk and Insurance*, 75(1): 83–100.

Browne, Mark J., Ellen S. Pryor, and Bob Puelz. 2004. "The Effect of Bad-Faith Laws on First-Party Insurance Claims Decisions." *The Journal of Legal Studies*, 33(2): 355–390.

Cabral, Luis, and Arthur Fishman. 2012. "Business as Usual: A Consumer Search Theory of Sticky Prices and Asymmetric Price Adjustment." *International Journal of Industrial Organization*, 30(4): 371–376.

Cabral, Luis, and Michael Riordan. 1994. "The Learning Curve, Market Dominance and Predatory Pricing." *Econometrica*, 62(5): 1115–1140.

Cabral, Luis, and Michael Riordan. 1997. "The Learning Curve, Predation, Antitrust, and Welfare." *The Journal of Industrial Economics*, 45(2): 155–169.

Cabral, Ricardo, and Michael J. Leiblein. 2001. "Adoption of a Process Innovation with Learning-by-Doing: Evidence from the Semiconductor Industry." *The Journal of Industrial Economics*, 49(3): 269–280.

Clark, Robert, and Jean-Francois Houde. 2013. "Collusion with Asymmetric Retailers: Evidence from a Gasoline Price-Fixing Case." *American Economic Journal: Microeconomics*, 5(3): 97–123.

Clark, Robert, and Jean-Francois Houde. 2014. "The Effect of Explicit Communication on Pricing: Evidence from the Collapse of a Gasoline Cartel." *The Journal of Industrial Economics*, 62(2): 191–228.

Compte, Olivier, Frederic Jenny, and Patrick Rey. 2002. "Capacity Constraints, Mergers and Collusion." *European Economic Review*, 46(1): 1–29.

Corp., General Reinsurance, and Edward Angell Palmer & Dodge LLP. 2008. "Bad Faith Laws for Property/Casualty Claims."

Corts, Kenneth S. 1999. "Conduct Parameters and the Measurement of Market Power." *Journal of Econometrics*, 88(2): 227–250.

Danzon, Patricia, and Lee A. Lillard. 1983. "Settlement out of Court: The Disposition of Medical Malpractice Claims." *The Journal of Legal Studies*, 12(2): 345–377.

Dasgupta, Partha, and Joseph Stiglitz. 1988. "Learning-by-Doing, Market Structure and Industrial and Trade Policies." *Oxford Economics Papers*, 40(2): 246–268.

Davidson, Carl, and Raymond Deneckere. 1990. "Excess Capacity and Collusion." *International Economic Review*, 31(3): 521–541.

Dick, Andrew. 1991. "Learning by Doing and Dumping in the Semiconductor Industry." *Journal of Law and Economics*, 34(1): 133–159.

Dionne, Georges, and Pierre St-Michel. 1991. "Worker's Compensation and Moral Hazard." *The Review of Economic Studies*, 73(2): 236–244.

Eaton, Thomas A., David B. Mustard, and Susette M. Talarico. 2005. "The Effects of Seeking Punitive Damages on the Processing of Tort Claims." *Journal of Legal Studies*, 34(2): 343–369.

Eisenberg, Theodore, and Henry Farber. 2003. "The Government as Litigant: Further Tests of the Case Selection Model." *American Law and Economics Review*, 5(1): 94–133.

Engle, Robert, and Clive Granger. 1987. "Co-Integration and Error Correction: Representation, Estimation, and Testing." *Econometrica*, 55(2): 251–276.

Farber, Henry S., and Michelle J. White. 1991. "Medical Malpractice: An Empirical Examination of the Litigation Process." *Rand Journal of Economics*, 22(2): 199–217.

Farber, Henry S., and Michelle J. White. 1994. "A Comparison of Formal and Informal Dispute Resolution in Medical Malpractice." *Journal of Legal Studies*, 23(2): 777–806.

Feinman, Jay M. 2012. "The Law of Insurance Claim Practices: Beyond Bad Faith." *Insurance Practice Law Journal*, 47: 693–740.

Fershtman, Chaim, and Ariel Pakes. 2000. "A Dynamic Oligopoly with Collusion and Price Wars." *The Rand Journal of Economics*, 31(2): 207–236.

Flamm, Kenneth, and Peter C. Reiss. 1993. “Semiconductor Dependency and Strategic Trade Policy.” *Brookings Papers on Economic Activity, Microeconomics*, , (1): 249–333.

Fournier, Gary M., and Thomas W. Zuehlke. 1989. “Litigation and Settlement: An Empirical Approach.” *The Review of Economics and Statistics*, 71(2): 189–195.

Fudenberg, Drew, and Jean Tirole. 1983. “Learning-by-Doing and Market Performance.” *The Bell Journal of Economics*, 14(2): 522–530.

Gardete, Pedro. 2014. “Should Capital-Intensive Firms share Demand Information with Competitors?” Working Paper, Stanford Graduate School of Business.

Genesove, David, and Wallace P. Mullin. 2001. “Rules, Communication and Collusion: Narrative Evidence from the Sugar Institute Case.” *American Economic Review*, 91(3): 379–398.

Ghemawat, Pankaj, and Michael A. Spence. 1985. “Learning Curve Spillovers and Market Performance.” *The Quarterly Journal of Economics*, 100: 839–852.

Gould, John P. 1973. “The Economics of Legal Conflicts.” *The Journal of Legal Studies*, 2(2): 279–300.

Gruber, Harald. 1998. “Learning-by-Doing and Spillovers: Further Evidence from the Semiconductor Industry.” *Review of Industrial Organization*, 13(6): 697–711.

Haltiwanger, John, and Joseph E. Jr. Harrington. 1991. “The Impact of Cyclical Demand Movements on Collusive Behavior.” *The Rand Journal of Economics*, 22(1): 89–106.

Harrington, Joseph E. 2004. “Cartel Pricing Dynamics in the Presence of an Antitrust Authority.” *Rand Journal of Economics*, 35(4): 651–673.

Harrington, Joseph E. 2005. “Optimal Cartel Pricing in the Presence of an Antitrust Authority.” *International Economic Review*, 46(1): 145–169.

Harrington, Joseph E., and Joe Chen. 2006. “Cartel Pricing Dynamics wtih Cost Variability and Endogeneous Buyer Detection.” *International Journal of Industrial Organization*, 24(6): 1185–1212.

Hatch, Nile W., and David C. Mowery. 1998. “Process Innovation and Learning by Doing in Semiconductor Manufacturing.” *Management Science*, 44(11): 1461–1477.

Hughes, James W., and Edward A. Snyder. 1989. “Policy Analysis of Medical Malpractice Reforms: What Can we Learn from Claims Data?” *Journal of Business & Economic Statistics*, 7(4): 423–431.

Hughes, James W., and Edward A. Snyder. 1995. “Litigation and Settlement under the English and American Rules: Theory and Evidence.” *Journal of Law and Economics*, 38(1): 225–250.

Irwin, Douglas A., and Peter J. Klenow. 1994. “Learning-by-Doing and Spillovers in the Semiconductor Industry.” *Journal of Political Economy*, 102(6): 1200–1227.

Johnston, Jason Scott, and Joel Waldfogel. 2002. “Does Repeat Play Elicit Cooperation? Evidence From Federal Civil Litigation.” *The Journal of Legal Studies*, 31(1): 39–60.

Kang, Joonkyu. 2010. “A Study of the DRAM Industry.” Master’s diss. Massachusetts Institute of Technology.

Katz, Avery. 1990. “The Effect of Frivolous Lawsuits on the Settlement of Litigation.” *International Review of Law and Economics*, 10(1): 3–27.

Kessler, Daniel. 1996. “Institutional Causes of Delay in the Settlement of Legal Disputes.” *Journal of Law, Economics, and Organization*, 12(2): 432–460.

Kessler, Daniel P., and Daniel L. Rubinfeld. 2007. “Empirical Study of the Civil Justice System.” In *Handbook of Law and Economics*. Vol. 1, , ed. A. Mitchell Polinsky and Steven Shavell, 343–402. Elsevier: North-Holland.

Kessler, Daniel, Thomas Meites, and Geoffrey Miller. 1996. “Explnaining Deviations from the Fifty-Percent Rule: A Multimodal Approach to the Selection of Cases for Litigation.” *The Journal of Legal Studies*, 25(1): 233–259.

Landes, William M. 1971. “An Economic Analysis of the Courts.” *Journal of Law and Economics*, 14(1): 61–107.

Landes, William M., and Richard A. Posner. 1979. “An Economic Approach to Legal Procedure and Judicial Administration.” *The Journal of Legal Studies*, 8(2): 235–284.

Levenstein, Margaret. 1997. “Price Wars and the Stability of Collusion: A Study of the Pre-World War I Bromine Industry.” *The Journal of Industrial Economics*, 45(2): 117–137.

Levenstein, Margaret, and Valerie Suslow. 2006. “What Determines Cartel Success?” *Journal of Economic Literature*, 44(1): 43–95.

Lewis, Matthew S. 2011. “Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market.” *Journal of Economics and Management Strategy*, 20(2): 409–449.

Macher, Jeffrey T., and David C. Mowery. 2003. ““Managing” Learning by Doing: An Empirical Study in Semiconductor Manufacturing.” *The Journal of Product Innovation Management*, 20(5): 391–410.

Mariuzzo, Franco, and Patrick Paul Walsh. 2013. “Commodity Market Dynamics and the Joint Executive Committee, 1880-1886.” *Review of Economics and Statistics*, 95(5): 1722–1739.

Mookherjee, Dilip, and Debraj Ray. 1991. “Collusive Market Structure under Learning-by-Doing and Increasing Returns.” *Review of Economics Studies*, 58(5): 993–1009.

Nalebuff, Barry. 1987. "Credible Pre-trial Negotiation." *The Rand Journal of Economics*, 18(2): 198–210.

Noll, Roger. 2014. "The DRAM Antitrust Litigation." In *The Antitrust Revolution: Economics, Competition and Policy*. . One ed., , ed. John E. Kwoka Jr. and Lawrence J. White, 246–276. Oxford University Press.

Nye, William W. 1996. "Firm-specific Learning-by-Doing in Semiconductor Production: Some Evidence from the 1986 Trade Agreement." *Review of Industrial Organization*, 11(3): 383–394.

Ostrager, Barry R., and Thomas R. Newman. 2008. *Handbook on Insurance Coverage Disputes.* Vol. 2. Fourteenth ed., New York: Aspen Publishers.

Perloff, Jeffrey M., Daniel L. Rubinfeld, and Paul Ruud. 1996. "Antitrust Settlements and Trial Outcomes." *Review of Economics and Statistics*, 78(3): 401–409.

P'ng, Ivan P. 1983. "Strategic Behavior in Suit, Settlement and Trial." *Bell Journal of Economics*, 14(2): 539–550.

Porter, Robert, and Douglas J. Zona. 1993. "Detection of Bid Rigging in Procurement Auctions." *The Journal of Political Economy*, 101(3): 518–538.

Porter, Robert, and Douglas J. Zona. 1999. "Ohio School Milk Markets: an Analysis of Bidding." *The Rand Journal of Economics*, 30(2): 263–288.

Posner, Richard A. 1973. "An Economic Approach to Legal Procedure and Judicial Administration." *The Journal of Legal Studies*, 2(2): 399–458.

Prescott, J.J., and Kathryn E. Spier. 2015*a*. "A Comprehensive Theory of Civil Settlement." Working Paper.

Prescott, J.J., and Kathryn E. Spier. 2015*b*. "Tailored Suits: Contracting on Litigation." Working Paper.

Prescott, J.J., Kathryn E. Spier, and Albert Yoon. 2014. "Trial and Settlement: A Study of High-Low Agreements." *Journal of Law and Economics*, 57(3): 699–746.

Priest, George L. 1987. "Measuring Legal Change." *Journal of Law, Economics, & Organization*, 3(2): 193–225.

Priest, George L., and Benjamin Klein. 1984. "The Selection of Disputes for Litigation." *The Journal of Legal Studies*, 13(1): 1–55.

Reinganum, Jennifer F., and Louise L. Wilde. 1986. "Settlement, Litigation, and the Allocation of Litigation Costs." *The Rand Journal of Economics*, 17(4): 557–566.

Roller, Lars-Hendrik, and Frode Steen. 2006. "On the Workings of a Cartel: Evidence from the Norwegian Cement Industry." *The American Economic Review*, 96(1): 321–338.

Rotemberg, Julio J., and Garth Saloner. 1986. "A Supergame-Theoretic Model of Price Wars during Booms." *American Economic Review*, 76(3): 390–407.

Scott Morton, Fiona. 1997. "Entry and Predation: British Shipping Cartels 1879-1929." *Journal of Economics and Management Strategy*, 6(4): 679–724.

Shavell, Steven. 1982. "Suit, Settlement and Trial: A Theoretical Analysis under Alternative Methods for the Allocation of Legal Costs." *The Journal of Legal Studies*, 11(1): 55–82.

Siebert, Ralph. 2010. "Learning-by-Doing and Cannibalization Effects at Multi-Vintage Firms: Evidence from the Semiconductor Industry." *The B.E. Journal of Economic Analysis & Policy*, 10(1): Article 37.

Siebert, Ralph, and Christine Zulehner. 2013. "The Impact of Market Demand and Entry Costs on Market Structure." Working Paper, Purdue University Krannert School of Management.

Siegelman, Peter, and Joel Waldfogel. 1999. "Toward a Taxonomy of Disputes: New Evidence through the Prism of the Priest/Klein Model." *The Journal of Legal Studies*, 28(1): 101–130.

Siegelman, Peter, and John J. Donohue. 1995. "The Selection of Employment Discrimination Disputes for Litigation: Using Business Cycle Effects to Test the Priest-Klein Hypothesis." *The Journal of Legal Studies*, 24(2): 427–462.

Spence, Michael A. 1981. "The Learning Curve and Competition." *The Bell Journal of Economics*, 12(1): 49–70.

Spier, Katherine. 2007. "Litigation." In *Handbook of Law and Economics*. Vol. 1, , ed. A. Mitchell Polinsky and Steven Shavell, 259–342. Elsevier: North-Holland.

Spier, Kathryn E. 1992. "The Dynamics of Pretrial Negotiation." *Review of Economics Studies*, 59(1): 93–108.

Stempel, Jeffrey W. 2006. *Stempel on Insurance Contracts.* . Third ed., New York: Aspen Publishers.

Sykes, Alan O. 1994. "Judicial Limitations on the Discretion of Liability Insurers to Settle or Litigate: An Economic Critique." *Texas Law Review*, 72: 1345–1374.

Sykes, Alan O. 1996. "'Bad Faith' Breach of Contract by First-Party Insurers." *The Journal of Legal Studies*, 25(2): 405–433.

Tappata, Mariano. 2009. "Rockets and Feathers: Understanding Asymmetric Pricing." *The Rand Journal of Economics*, 40(4): 673–687.

Tennyson, Sharon, and William J. Warfel. 2010. "The Law and Economics of First-Party Insurance Bad Faith Liability." *Connecticut Insurance Law Journal*, 16(1): 203–242.

Verlinda, Jeremy A. 2008. "Do Rockets Rise Faster and Feathers Fall Slower in an Atmosphere of Local Market Power? Evidence from the Retail Gasoline Market." *Journal of Industrial Economics*, 26(6): 1290–1307.

Viscusi, W. Kip. 1988. "Product Liability Litigation with Risk Aversion." *The Journal of Legal Studies*, 17(1): 101–121.

Waldfogel, Joel. 1995. "The Selection Hypothesis and the Relationship between Trial and Plaintiff Victory." *The Journal of Political Economy*, 103(2): 229–260.

Waldfogel, Joel. 1998. "Reconciling Asymmetric Information and Divergent Expectations Theories of Litigation." *Journal of Law and Economics*, 41(2): 451–476.

Yang, Huanxing, and Lixin Ye. 2008. "Search with Learning: Understanding Asymmetric Price Adjustments." *The Rand Journal of Economics*, 39(2): 547–564.

Zulehner, Christine. 2003. "Testing Dynamic Oligopolistic Interaction: Evidence from the Semiconductor Industry." *The International Journal of Industrial Organization*, 21(10): 1527–1556.