The Value of Market Information in the Dynamics of a Capital-Intensive Industry: The Case of DRAM Manufacturing*

Pedro Gardete

(Stanford University)

January 2012

Abstract

Much like in other semiconductor environments, DRAM manufacturers face significant demand uncertainty before production and capacity decisions can be implemented. This paper investigates the role and assesses the value of market information in the DRAM manufacturing industry. A dynamic oligopoly model of competition with capacity constraints is developed, where firms hold information about demand. The model reveals that better information is especially beneficial when the market performs poorly. In addition, a non-linear relationship between the quality of the information and the performance of the industry is found, once the costs of adjusting capacity levels are incorporated. The equilibrium is solved through an approximation of the firms' actions conditional on the market information available to them, and the parameters are estimated through the simulated method of moments. The model recovers the market information that rationalizes the decisions of the firms. The available market information is valued at approximately 290 million dollars for the industry per quarter, and is estimated to have yielded returns of 1,6 billion dollars during the recent DRAM market 'bust' period, from 2008 to mid 2009.

*I am indebted to Miguel Villas-Boas and Denis Nekipelov for their guidance and valuable advice. I am also indebted to Ron Berman, Debbie Chiou, Chan Jean Lee, TI Kim, John Morgan, Ram Rao, Diogo Silva, Shubhranshu Singh and Kaifu Zhang for their valuable suggestions.

Correspondence: Stanford University, Graduate School of Business, 655 Knight Way, Stanford CA, 94305, USA. Email: gardete@stanford.edu.

1 Introduction

Consider the following sequence of events: on August 15th 2007, the EETimes reported that the DRAM (Dynamic Random Access Memory) industry had reasons "for cautious optimism," mostly because of signs of stronger demand ahead driven by sales of handsets and game consoles. Two weeks later, iSuppli revised its own market forecast with a $2.5 billion cut in expected annual revenues, but announced a prediction of a 17.5% growth in sales for 2008. By mid-2008, the president of the Semiconductor Industry Association, George Scalise, reported that DRAM revenues had declined by 34% in the first four months of 2008.[1]

The DRAM industry is only one of many industries that depend heavily on assessments of the economic context to make a profit. Competitors face a number of decisions - including production orders and capacity investments - that are made before a significant amount of uncertainty is resolved. In this industry the lag that spans from deciding capacity investments to their full implementation is of at least one quarter, which need an additional quarter for the production units to go through the production process and be brought to the market. These lag times in conjunction with high demand uncertainty create a turbulent setting for all the firms in the DRAM industry.[2]

This paper analyzes how the industry evolution depends on the information held by firms, in a competition setting. It takes into account the fact that managers rely on market information to make decisions, and that they expect their competitors to do the same.

Because accessing reliable and complete data on public and private information is infeasible, I instead recover the *unobservable* information that is consistent with the *observable* data on firms' actions by use of a game-theoretic model. The model used for analysis contemplates decisions that carry present and future consequences. Production decisions affect current profit levels according to the true realization of demand, while capacity decisions carry dynamic effects since they determine the competitive conditions for the future.

After estimating an econometric model from first-principles it is possible to construct different scenarios and describe how the industry is affected by them. These scenarios allow us to calculate

[1] Sources: www.marketwatch.com/story/pressures-mounting-anew-on-dram-chip-market-isuppli-says, www.eetimes.com/showArticle.jhtml?articleID=201800294, and www.goo.gl/QWTh.

[2] The October 19th 2009 edition of the Wall Street Journal features a piece describing the volatility of the DRAM industry with an illustrative title: "The Chip Sector's Boom or Bust DRAMa".

the precision and value of the information in the industry.

The DRAM market is presented next. Section 3 reviews the related literature. The dataset used for the analysis is described in Section 4. Section 5 presents the theoretical model used for the analysis and Section 6 describes the model implementation and estimation. Section 7 presents and discusses the results and Section 8 describes counterfactual analyses that are used to value market information. Section 9 concludes.

2 Description of the DRAM Market

DRAM is a technological variant of what is generically called RAM or 'computer memory'. It is used in a number of electronic appliances like computers, cell phones, printers, game consoles among others. Intel is generally credited for having introduced the first DRAM chip in 1970: the 1103 model featuring 1Kb of storage capacity - 1024 bits - which competed with the dominant technology until then, the magnetic core. Besides Intel, two other American companies started DRAM production in the same year with similar products: Mostek and Advanced Memory Systems (the latter effectively entering the market before Intel, but with a less popular product).[3] The success of DRAM over its predecessor technology was primarily due to its compactness which translated into less space required for each bit of storage, and in addition due to its faster read/write performance. DRAM dominated the volatile memory industry in less than a decade. A number of fundamental characteristics of the market were soon apparent and still survive until today. They are discussed below in turn.

2.1 A Race for Memory Density

The production of DRAM chips requires hundreds of procedural steps, and has been described as "one of the most difficult processes of high technology known to man" (Murillo, 1993). This process has only gotten more complex over time, as there has been a continued effort to introduce more memory into the same amount of physical space.

Memory chips are etched on a silicon disc called wafer (see Figure A in the appendix for a picture of one). Given that the cost of a wafer is basically constant across production technolo-

[3] See Murillo (1993) for a description of the DRAM industry from 1970 until 1993.

gies, it follows that fitting double as many memory units into the same wafer results in halving the variable production costs. In fact, memory density is the fundamental cost determinant for the DRAM production process. Production decisions depend strongly on it, over and above other cost factors such as labor wages or the price of silicon and of other raw materials.

Production yield - the usable number of chips per silicon wafer - is a key success factor in the industry. In fact, the 1103 chip's early market leadership occurred partly because of Intel's achievement of a 10% die yield, while at the same time competitors were known to achieve yields of 3% to 6%. As of 1997, a UC Berkeley study showed that the typical average line yield (the percentage of wafers that were not scrapped from the production process) was of 93%, and the die yield was of 77.4%, resulting in a ~72% effective chip yield for 200mm wafers (see ICE Corporation, 1997, Chp. 3).[4]

2.2 International Competition

Except for its earlier stages, competition in the DRAM industry has taken place at the international scale. The introduction of the 1103 chip in 1970 had already caused the appearance of clones by other 19 companies as of 1972, including the first Japanese rival, NEC. By 1976 two Japanese companies had subsidiaries in the U.S.A., and one of them (NEC) was receiving its first big order of DRAM chips from Honeywell ($5 million worth). This was the fruit of continued efforts by the Japanese companies to address reliability and provide fast access performance in their chips. Soon, the so-called 'four sisters' (NEC, Toshiba, Fujitsu and Hitachi) had managed to effectively enter the American market.[5]

In October 1985 Intel announced it was leaving the DRAM production altogether in order to focus on the microprocessor, a bet that proved successful *a posteriori*. In the same year Motorola declared it too was abandoning the market.[6] A period of Japanese dominance in the DRAM market had started, even if short-lived: in September 1986, a U.S.A.-Japan Semiconductor Agreement forced Japanese firms to stop "dumping" across world markets, and established

[4]New generations of DRAM memory are launched over time, substituting previous ones gradually (see for example Norton and Bass (1987)). In this paper I consider the smooth average effect of such successive generations.

[5]A number of enabling factors are documented: availability of qualified human resources, low cost of capital, "deep pockets" linked to the Keiretsu firm structure, among others. By 1980 the differences in product reliability were measured repeatedly, giving significant advantage to the Japanese producers.

[6]Both firms reentered the market in 1987 as resellers.

"fair market values" for their products. The agreement was aimed at benefiting American and European producers through raising Japanese product prices. Instead, it changed the competitive landscape in an unpredicted fashion: it opened the door for Korean firms who had been acquiring technology licenses and had started to participate in the capital of American companies (e.g.: Samsung purchased 2.7% capital of the American DRAM manufacturer Micron in the mid-80's). Concurrently, the quality-orientation that had served Japanese firms well in the past may have played a role in their demise. The shift in use of DRAM from mainframes to personal computers privileged the quest for low costs instead of quality and reliability.[7] Korean firms eventually surpassed both Japanese and North American ones in production quantities, as described in Table 1. Currently, the DRAM market is organized into roughly more than 10 players, many of which are organized into alliances of firms which effectively share production technologies and coordinate R&D efforts, decide production levels and allocate available capacity across its members.

2.3 Turbulent Prices Descents

Price descents have been part of the DRAM market since its beginning. For instance, the price of the 1103 chip came down from $60 in 1970 to only $4 three years later. The downward trend in prices per bit has remained until today. The reference spot price for a Kilobit of DRAM was of 0.0003 cents of a dollar as of April 2010,[8] a reduction of 7 orders of magnitude in price per memory unit with respect to the original level of 1970. The improvements in production technologies are the main factor responsible for long-term decreases in market prices. Similar to what happens in other semiconductor industries, the price per unit decline is a result of firm coordination around Moore's law. In the DRAM industry, this law has been translated to state that the number of bits per wafer doubles every two years.[9] Moore's law is today an effective focal point that firms use to coordinate the rhythm of increases in production efficiency, much in the spirit of Schelling (1960). It effectively drives prices down for semiconductor components over time.

Another empirical regularity is that the price descent path is not smooth. Economic condi-

[7] As an example, today's DRAM modules used in computer servers often have error checking and error-correcting mechanisms (e.g.: parity bits), which are absent in most of personal computers due to cost reasons.

[8] As noted in www.dramexchange.com on April 13th, 2010.

[9] See intel.com/pressroom/kits/events/moores_law_40th/.

tions impact the performance of the DRAM industry, pushing prices up and down around the descent path over time.

3 Related Literature

There exist a number of theoretical contributions analyzing the role of information in market competition. Some examples are Novshek and Sonnenschein (1982), Clarke (1983), Vives (1984), Gal-Or (1985; 1986) and Shapiro (1986). They analyze the roles of information and the incentives of firms to share it in market settings.[10]

The most prevalent finding is that better information has ambiguous effects on the profitability of firms, depending on the market structure (see Raith (1996); Jin (2000)). For example, under a linear demand specification and constant marginal costs, sharing information about demand is found to be optimal under Bertrand competition but not under Cournot, while sharing information about costs is found to be optimal under Cournot competition, but not under Bertrand. Traditionally, the effects of information sharing have been explained by how it affects the correlation between the actions of the firms and whether those actions are strategic complements or substitutes. For example, sharing news about market demand will increase the correlation of production quantities. Because quantities are strategic substitutes under the linear demand specification, the net effect of information sharing is negative. On the other hand, when a firm shares private information about its low costs, it increases its production quantities and decreases the production of its rival. When firms face a linear demand, the negative correlation of quantities is beneficial for the firms since best-response functions are negatively sloped. The dual argument may be made for Bertrand competition. However, no prediction is available once capacity constraints are introduced, or once different demand specifications are considered.

The analysis of the role of information and its characteristics in strategic settings has been an almost exclusive domain of research in auctions. A noteworthy exception is provided by Armantier and Richard (2003) who analyze the impact of sharing cost information in the airline industry. They consider a static duopoly model where carriers can decide whether to operate in several markets and how much to 'produce' in each one simultaneously. Sharing cost information

[10] The role of market information in competition contexts has also been analyzed in the context of distribution channels (e.g.: Chu, 1992; Lariviere and Padmanabhan, 1997) and of specific settings of information sharing (e.g.: Villas-Boas, 1994; Chen, Narasimhan, and Zhang, 2001).

significantly increases firms' profits while it only moderately decreases consumer surplus overall. The focus of the present work is on analyzing the role of information about demand conditions instead of about production costs. In addition this paper looks at dynamic effects of market information, and by incorporating computational advances is able to do so in non-linear demand settings.

While there exist several empirical contributions in dynamic market competition settings that include private information (see for example Pesendorfer and Schmidt-Dengler, 2003; Bajari, Benkard, and Levin, 2007; Aguirregabiria and Mira, 2007; Bajari, Chernozhukov, Nekipelov, and Hong, 2009; Ryan, 2009; Fershtman and Pakes, 2009), most of their attention is devoted to advancing estimation methods and none focuses on the role of information in markets. In addition, while these models include uncertainty in the cost structure, the DRAM industry (and semiconductor industries in general) provides a contrasting setting where the cost structure is predictable and fairly public, while demand is highly uncertain and plays an important role.

4 Data

The dataset has quarterly information from the first quarter of 2005 until the third quarter of 2010. The period of analysis takes place after the complaints by US personal computer makers of collusive behavior by DRAM makers. These claims were proven to be true: they led to heavy penalties to DRAM companies in the years of 2003 to 2005.[11] The executives of some of the most important firms in the industry - Samsung, Hynix, Elpida and Infineon - pleaded guilty to either price fixing or obstructing justice in relation to a United States Justice Department's probe on the matter.

A firm in the DRAM industry provided the proprietary data for the model estimation. The data comprises detailed price, production, capacity and cost information on seven of the main firms in the DRAM industry who account for roughly 90% of the market share of the global DRAM market. These firms are organized into four alliances which effectively behave as coordinated organisms (and hence are taken as individual decision-making units in the analysis). Table 2 presents the descriptive statistics of the dataset. In the beginning of the sample and until the third quarter of 2008, 'virtual' capacity usage (not taking outages and maintenance

[11] I assume that by the start of 2005 these firms already had incentives to drop any collusive behaviors.

stoppages into account) was at 100% and the DRAM market faced undersupply. The main problem for firms was to update their production capacity appropriately so as to be able to sell more units at lower costs. However, from the last quarter of 2008 until late 2009 the sample period capacity utilization dropped dramatically. During that period marginal production costs played an important role in the production levels of firms. Simultaneously, firms refinanced heavily through either their parent companies or through renegotiation with banks in order to stay in operation. The medium-sized player Qimonda (not present in the sample data) was the single casualty after having filed for bankruptcy in the beginning of 2009. A listing of the firms included in the analysis as well as of their nationalities and partnerships is provided in Table 3.[12]

While producing DRAM implies putting a different number of resources in place, it is by far the production technology that determines the relevant cost for decision-making about production. The cost of producing a 'wafer' is relatively constant across technologies. However, the memory density/efficiency of a wafer (i.e.: the number of bits per wafer) varies by production technology and ultimately dictates the relevant production cost per unit (measured in millions of gigabytes/GB). These data are not published and are only accessible through expert analysis. For the purposes of this paper several meetings with marketing managers of a firm in the industry provided information on the density levels for the alliances in the sample.

South Korean Samsung Electronics is the leader in market share throughout the sample period, capturing almost a third of the unit sales in the market. It is followed by Hynix which achieves more than 20% market share of DRAM unit sales. The remaining chunks of the market are shared among the medium size competitors (i.e.: Elpida and Micron with 19% and 13% market shares), and by smaller firms (Powerchip and ProMOS with less than 6% each).[13]

While production increased for most firms along the sample period (and all firms were producing more in mid-2009 than in the beginning of 2005), capacity usage dropped from full utilization (which occurs until the third quarter of 2008) to roughly 80% until the third quarter of 2009. Finally, market price and production costs follow a downward trend over time.

The discrepancy between capacity levels and unit productions depends on both capacity

[12] Elpida joined Micron in an alliance near the end of the sample. This is a unique event in the data and is incorporated exogenously in the analysis.

[13] See for example http://www.isuppli.com/Memory-and-Storage/News/Pages/Taiwans-Powerchip-Surges-in-First-Quarter-DRAM-Ranking.aspx.

and production decisions by firms. These in turn, depend on two sets of factors: capacity and production costs on the one hand, and on market information on the other. It is therefore important to incorporate both these sets of factors into the game theoretic model.

5 The Model

This section presents a model of competition for a capital-intensive industry. The main decisions of the firms are setting their production and capacity levels. At a given period t, firms decide how much production and capacity they want to have available in the next period. Hence, at time t they decide production and capacity levels for period $t+1$. In that next period, firms sell their production at the market price, earn economic profits/losses and go back to making further decisions. These decisions are dependent on each other: at period t the decision on production levels for the next period is bounded above by the available production capacity, which was chosen in the previous period. Hence, capacity levels only have a market impact two periods after they have been decided upon.

A key component of the model is that firms do not have perfect information about the market performance but instead use market information to guide their decisions. The structure of the information is as follows: at time t firms hold private signals about the demand performance on time $t+2$. These demand signals are correlated, since they all inform the firms about the future demand conditions (albeit the correlation can be arbitrarily close to zero). In order to incorporate the idea that firms converge on the available information over time, the private demand signals become public after one period elapses. There is a consequence of this assumption to the decision-making process: because capacity decisions are taken considering the market outlook two periods from today, they use only private information. However, because quantity decisions need only consider the market conditions tomorrow, they can be informed by the demand signals about the next period's performance, which used to be private in the previous period but have become public today, after one period has elapsed. Hence, firms play a private information game on capacity levels, and a public information game on production levels.

This information structure has two key advantages. First, it captures the idea that firms learn more as they get closer to the market interaction stage, and that market forecasts become common knowledge in the industry over time. Second, it increases the tractability of the model.

The timing of the game is as follows:

Figure 1: Timing

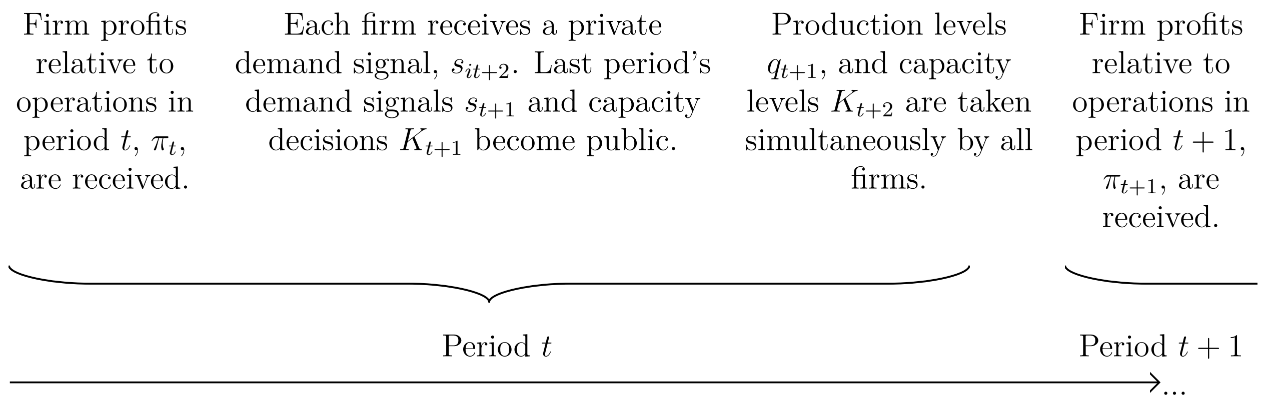

The timing above uses notation that is explained now: q_{it+1} denotes the production decision of firm i that will be brought to the market at period $t+1$. Decision K_{it+2} denotes the capacity level used for the production that will reach the market in period t+2. Because of the lag between decisions and outcomes, at time t firm i makes decisions q_{it+1} and K_{it+2}. Hence, at time t the firm decides how much to produce for the next period, q_{it+1}, constrained by the capacity decision it made in the previous period, K_{it+1}. Finally, at time t, s_{it+2} denotes the private information that firm i holds about the market conditions at time $t+2$, and s_{t+1} represents a vector of demand signals about the market conditions at time $t+1$. These signals used to be private two periods ahead, in period $t-1$, but have become public in period t.

5.1 Demand and Information Structure

A demand curve for a homogeneous good is considered, yielding the inverse demand function:

$$P_t = P\left(Q_t, \theta, \varepsilon_t\right), \, \varepsilon_t \sim F_\varepsilon\left(\cdot\right)$$

where P_t is the market price for period t, Q_t is the aggregate production at period t, θ is a set of demand parameters and ε_t is a demand shock that influences the market willingness-to-pay for each unit of the good and is unknown to the firms *a priori*. The demand shock ε_t reflects the total market uncertainty under absence of information. While firms do not observe it directly, they hold noisy information about the demand shock in the form of a signal $s_{it} = f(\varepsilon_t, \eta_{it})$ with $\eta_{it} \sim F_\eta(\cdot)$, such that it is centered at the true demand shock, i.e. $E[s_{it}] = E[\varepsilon_t]$, $i = 1, .., n$.[14] The demand signals summarize the firms' market information. In reality it can be comprised of different sources of information, from market reports, managerial intuition, market research, etc.

5.2 Firm Decisions

5.2.1 Production Decisions

Firms decide their production levels before knowing the true demand realizations and consequently before knowing the market price they will face. At the production stage they hold common information about the expected market performance in the next period. Hence, at time t firm i decides its production level for the next period by solving the objective function

$$\begin{aligned} \underset{q_{it+1}}{max} \quad & E_{\varepsilon_{t+1}}\left[P\left(Q_{t+1}, \theta, \varepsilon_{t+1}\right) - c_{it+1} \middle| s_{t+1}\right] q_{it+1} \\ s.t. \quad & q_{it+1} \leq K_{it+1} \end{aligned} \tag{1}$$

where c_{it+1} denotes firm i's marginal cost, and $s_{t+1} = [s_{1t+1} \cdots s_{nt+1}]'$ holds the short-run public signals, which used to be private information of each firm at time $t-1$ but have become public at time t. Notice that because K_{it+1} was decided in period $t-1$, it is set in period t and used as given.

5.2.2 Capacity Decisions

Capacity levels take one period to be set in place and made available for production. Since there exists an additional production lag of one period, capacity levels affect market outcomes two periods after their decision was originally made. Because they affect future outcomes, firms should be forward-looking while setting them. At time $t-2$ firms decide their capacity level taking future production costs and private information signals into account, as well as other state variables such as the current installed capacity levels. Let $\Omega_t = \{c_{t+1}, c_{t+2}, K_{t+1}\}$ be a

[14] The letter n is used to denote the number of firms or the set of competing firms throughout the paper, depending on the context.

set of state variables, where c_{t+1} and c_{t+2} are vectors of the firms' marginal costs at times $t+1$ and $t+2$ respectively. In addition, let $\Theta_{it} = \{s_{it+2}, s_{t+1}\}$ be the set of information-related state variables, where s_{it+2} is a private information signal about demand in period $t+2$, and s_{t+1} is a collection of publicly-available demand signals about demand in period $t+1$. At time t firm i solves the dynamic problem that can be expressed as in the Bellman equation below:

$$\begin{aligned} V_i(\Omega_t, \Theta_{it}) = \delta \underset{q_{it+1}}{max} E\left[\pi_{it+1} | \Omega_t, \Theta_{it}\right] \; &+ \; \underset{K_{it+2}}{max} \, \delta E_{\varepsilon_{t+2}, s_{-it+2}}\left[V_i(\Omega_{t+1}, \Theta_{it+1}) | \Omega_t, \Theta_{it}\right] \qquad (2) \\ &- \; InvCost\left(K_{it+1}, K_{it+2}\right) \\ & \quad s.t. \, q_{it+1} \leq K_{it+1} \qquad (3) \end{aligned}$$

where the first component captures the production decision of the firm which will influence the expected profit in period $t+1$, and the second component captures the subsequent expected future profits, which will be influenced by the capacity decision, K_{it+2}. Finally, function $InvCost\left(K_{it+1}, K_{it+2}\right)$ captures the investment cost of changing capacity levels from K_{it+1} to K_{it+2}. Two remarks are in order: first, the profit in period t is sunk, and hence is ignored in the Bellman equation and in the subsequent decisions. Second, solving for the equilibrium capacity levels is more complex than solving for the static production decisions, because of two main reasons. First, capacity decisions impact the firms' competitive conditions into the future, so firms need to be forward-looking while taking them. Second, capacity decisions are taken under correlated private information. This means that a firm's private demand signal is informative not only about the future market conditions but also about its rivals' demand signals. It follows that firms must take into account the potential equilibrium actions the other firms may take depending on the information they are likely to hold.

5.2.3 Equilibrium Concept

I focus on Markov Bayesian perfect equilibria (see Maskin and Tirole, 2001), where only payoff-relevant state variables are used for decision-making, and firms take into account the probabilities of facing different types of competitors (based on their own type). In order to see this in Bellman equation (2), consider the case of competition between two firms. The second term of the right-hand side of the Bellman equation is expanded and presented below:

$$\underset{K_{it+2}}{max}\, \delta E_{\varepsilon_{t+2}, s_{-it+2}} \left[V_i \left(c_{t+2},\, c_{t+3},\, K_{it+2}, K^*_{jt+2}\left(s_{jt+2}\right)\right) \middle| \, c_{t+1},\, c_{t+2},\, K_{t+1}, s_{it+2}, s_{t+1} \right]$$

Firm i takes expectations over firm j's private demand signal, which becomes relevant through the equilibrium response $K^*_{jt+2}\left(s_{jt+2}\right)$. Hence, solving the equilibrium actions for firm i implies knowledge of the equilibrium policy of firm j, $K^*_{jt+2}\left(s_{jt+2}\right)$. Alternatively, we are looking for the equilibrium policies $K^*_{it+2}\left(s_{it+2}\right)\ \forall i \in n$ that satisfy a set of first-order conditions. Hence, the Bellman equation is a *functional* equation with respect to the value function as well as with respect to the equilibrium capacity policies. An equivalent interpretation is that we want to find capacity level K^*_{it+2} for each of the (most likely infinite) values of the demand signal, s_{it+2}.

6 Estimation

This section presents the implementation of the model within the context of the DRAM industry and outlines the estimation strategy.

6.1 Demand for DRAM and Information Structure

DRAM is a significantly commoditized product and is sought by electronic product manufacturers. Firms are assumed to face a constant-elastic inverse demand curve for a homogeneous good,

$$P_t = \alpha\, Q_t^{\beta}\, \varepsilon_t,\, \varepsilon_t \sim LogN\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right) \tag{4}$$

where ε_t denotes the stochastic component of price. The constant elastic specification arises as a result of a more general demand model nesting alternative demand specifications. It is also consistent with the industry's perception that the computer market is able to absorb increasing amounts of DRAM units at positive prices. The use of a linear specification would instead assume that a "choking quantity" exists such that market price would eventually be equal to zero.

The assumption that the demand shock ε_t is lognormally distributed has two main advantages: first, it restricts the willingness to pay of the market to the positive domain. Second, it ensures market prices have a constant variability in relative terms. This is especially important in markets such as the DRAM industry, where prices exhibit strong trends. In these cases the multiplicative specification avoids the stochastic component from subsiding, or alternatively, from overwhelming the deterministic component over time.[15]

Finally, the parametrization of the error distribution ensures that its mean is equal to one (this normalization is analogous to the zero mean assumption in additive models) and its variance is equal to $\exp\{\sigma_\varepsilon^2 - 1\}$. The demand shocks are interpreted in a percent scale, i.e., $\varepsilon_t = 1.1$ means that the willingness-to-pay of the market was 10% above its mean, given the aggregate market output. They represent other factors that affect demand, such as economic conditions, shocks to producers/consumers in substitute and/or complement product-markets, etc.

Demand signals are assumed to be uncorrelated over time and to relate to the contemporaneous demand shock as:

$$s_{it} = \varepsilon_t.\eta_{it},\ \eta_{it} \sim LogN\left(-\frac{\sigma_\eta^2}{2}, \sigma_\eta^2\right) \tag{5}$$

such that the mean of each of the components of the demand signals is equal to 1.[16] A more practical specification for estimation is attained after taking logarithms of both sides of (4):

$$\log\left(P_t\right) = \alpha^{'} + \beta \log\left(Q_t\right) + u_t,\ u_t \sim N\left(0, \sigma_\varepsilon^2\right) \tag{6}$$

where $\alpha^{'} \equiv \log(\alpha) - \frac{\sigma_\varepsilon^2}{2}$ and $u_t \equiv \log(\varepsilon_t) + \frac{\sigma_\varepsilon^2}{2}$.

Given the assumptions above one can recover convenient posterior probabilities. For example, in the case of two firms one gets the useful conditional distributions:

$$\begin{array}{c} s_{1t} \\ s_{2t} \end{array} \Bigg| \varepsilon_t \sim LogN\left(\begin{bmatrix} log\left(\varepsilon_t\right) - \frac{\sigma_\eta^2}{2} \\ log\left(\varepsilon_t\right) - \frac{\sigma_\eta^2}{2} \end{bmatrix}, \begin{bmatrix} \sigma_\eta^2 & 0 \\ 0 & \sigma_\eta^2 \end{bmatrix}\right)$$

[15] In addition, the fact that the observed prices follow a strong downward trend would suggest including a trend term in the demand specification. Such inclusion raised two issues: first, it explained away the production levels. Second, the negative trend coefficient found in that estimation would mean that the willingness-to-pay for the same quantity of bits has decreased over time, which has no industry validation.

[16] The model was also estimated with serially correlated demand shocks: it became challenging to disentangle the serial correlation component from the parameter capturing information precision, σ_η^2. Doing so would require a larger dataset. This does not affect firms' equilibrium behavior significantly: The model implementation in Section 6.3 reveals that firms need only anticipate the market two periods ahead, for which they have informative demand signals already available.

and

$$\varepsilon_t \left| s_{1t}, s_{2t} \sim LogN \left(\sigma_\varepsilon^2 \frac{\left(log\left(s_{1t}\right) + log\left(s_{2t}\right) + \sigma_\eta^2 \right)}{2\left(2\sigma_\varepsilon^2 + \sigma_\eta^2\right)}, \frac{\sigma_\varepsilon^2 \sigma_\eta^2}{2\sigma_\varepsilon^2 + \sigma_\eta^2} \right) \right. .$$

Because firms have access to information about market shocks through the demand signals, it is important to correct for endogeneity. Although the use of input prices is a popular choice, they are mostly irrelevant in the DRAM industry, as discussed in the introduction. The main production factor is defined by the density technology in place, on which managers base their production and capacity decisions strongly. Since memory density levels depend strongly on firms following Moore's law but are not related to demand shocks and demand signals, they are valid instruments for the production quantities. The matrix of instruments Z is formed by a constant and the memory densities across firms. These are used to identify the demand parameters under the "orthogonality" between the instrument matrix and the shifted logarithm of the error term, $E\left[Z'.u\right] = 0$. The demand parameters $\alpha, \beta, \sigma_\varepsilon^2$ were estimated jointly with the static parameters ω_0, σ_η^2 as described in the next section. The 2-step generalized method of moments was used, where $\widehat{\sigma_\varepsilon^2}$ is attained by first calculating the sample variance of the residuals and then inverting the variance of the log-normal distribution, i.e., $\widehat{\sigma_\varepsilon^2} = log\left(1 + \widehat{V\left(\varepsilon_t\right)}\right)$.

6.2 Firm quantity decisions

Firms decide how much to produce one period ahead of taking DRAM units into the market. Their production decisions are constrained above by their capacity level and in addition are affected by their demand information and cost structure. Their capacity decisions are made one period in advance, and thus are fixed during the subsequent production decisions. Besides capacity, production costs and availability of market information are the other components affecting production decisions. Given the demand and information structures above, the production problem of firm i can be written as

$$\begin{array}{ll} \underset{q_{it+1}}{max} & E_{\varepsilon_{t+1}}\left[\alpha \left(\sum_{j\in n} q_{jt+1} \right)^\beta \varepsilon_{t+1} - c_{it+1} \middle| s_{t+1} \right] q_{it+1} \\ s.t. & q_{it+1} \leq K_{it+1} \end{array} \tag{7}$$

where firm i's marginal production cost is constant and known, and given by

$$\underbrace{c_{it}}_{\$/GB} = \underbrace{\omega_0}_{Cost\, per\, Wafer (\$/Wafer)} \times \underbrace{\tau_{it}}_{\%\, Area\, per\, GB_{it} (Wafer/GB)}$$

where ω_0 is the cost of acquiring and processing a silicon wafer - which remains fairly constant over time and across firms - and can be estimated from the data. The scalar τ_{it} is a measure of production (in)efficiency. It measures the percentage area of a wafer that is used for each GB produced and is sourced from the data. As expected, $\{\tau_{it}\}$ decreases over time and is different across firms. Figure 2 presents the equilibrium production levels for four firms given an expected market performance $E(\varepsilon_{t+1})$. The equilibrium quantities are increasing on the expected demand shock up to reaching the installed capacity levels.

The production-related parameters ω_0 (wafer cost) and σ_η^2 ('variance' of the market information) are estimated through solving the quantity equilibrium described above for each time period in the data, and matching the moments of the data with those predicted by the model. In particular I match the first four moments of the distributions: $E\left[Q_t^l - \widehat{Q}_t^{*l}\right]$, $l = 1..4$. The first two moments intuitively identify the wafer cost and the variance of the demand signals. This is because the wafer cost affects mainly the mean level of production, and the quality of the demand information affects mainly the variance of the production. The latter moments are added so as to increase the efficiency of the estimation. Finally, $\widehat{Q}_t^*$ is defined as the predicted production by the model, i.e., $\widehat{Q}_t^* \equiv E\left(\sum_{i \in n} q_{it}^*(s_{it}) \middle| \widehat{\varepsilon}_t\right)$, such that it is the sum of the equilibrium quantities at an estimation guess of the cost and of the information precision parameters, given that the demand signals for the period are 'centered' at the demand shock that is recovered from the guesses of the demand curve parameters. Hence, for the purpose of estimation, firms get information centered at the recovered demand shocks, but do not have access to those shocks directly. The fact that the demand and the firm-side parameters are jointly estimated improves the efficiency of the estimator.

Predicting the production decisions entails solving the system of Kuhn Tucker conditions for each firm $i \in n$ at time t:

$$\begin{aligned} \alpha Q_t^{\beta-1}\left(Q_t+q_{it}\right) E_{\varepsilon_t}\left[\left.\varepsilon_t\right| s_t\right]-c_{it}-\lambda_{it} &= 0 \\ \lambda_{it}\left(K_{it}-q_{it}\right) &= 0 \\ \lambda_{it} &\geq 0 \\ q_{it} &\leq K_{it} \end{aligned}$$

An intuitive way of solving the system of Kuhn-Tucker conditions above while incorporating the demand shocks is to first generate demand signals from the density $f\left(\left.s_t\right| \widehat{\varepsilon}_t\right)$, and then to solve the quantity equilibrium at each set of simulations by calculating $E_{\varepsilon_t}\left[\left.\varepsilon_t\right| s_t\right]$ at each set of simulated signals.[17] The market output $\widehat{Q}_t^*$ can then be predicted by summing the expected individual production levels for the period. This process mimics the generation of some demand information s_t that firms are likely to have received according to the period's demand shock, given the guess of the parameters σ_{ε}^2 and σ_{η}^2. Two improvements are added so as to increase the speed of estimation to a practical amount of time. First, the density function of the vector of demand signals $f\left(\left.s_t\right| \widehat{\varepsilon}_t\right)$ is but a function of the sum of the demand signals, i.e., $f\left(\left.s_t\right| \widehat{\varepsilon}_t\right)=f_1\left(\left.\sum_{i \in n} s_{it}\right| \widehat{\varepsilon}_t\right)$. This is due to the distributional assumptions on the demand signals as well as to the fact that at the production stage the demand signals are public information. Second, the use of a Gauss-Hermite quadrature can drastically decrease the number of times the Kuhn-Tucker conditions are solved in order to find $\widehat{Q}_t^*$. The procedure is as follows: for each observation and guess of the demand parameters, I recover the estimated demand shocks $\widehat{\varepsilon}_t=\frac{P_t}{\widehat{\alpha} Q_t^{\widehat{\beta}}}$ and corresponding parameter estimate $\widehat{\sigma_{\varepsilon}^2}$. I then construct a Gauss-Hermite quadrature $G_1\left(\left.\sum_{i \in n} s_{it}\right| \widehat{\varepsilon}_t\right)$ and calculate the quantity equilibrium at each point of quadrature G_1 taking the capacity levels observed from the data as given. G_1 allows me to calculate the predicted market production for each period.[18] By applying a weighted sum one then gets the expected production quantity q_{it}^* for each firm that is then summed to those of its rivals in order to get the market expected period's production $\widehat{Q}_t^*$. The speed attained with this method is

[17] For the purpose of the speed of estimation, the algorithm tried successive candidates for satisfying the Kuhn-Tucker conditions and stopped when it found one. However, many experiments were done over different values of the parameters and not a single time did the procedure find more than a single equilibrium configuration.

[18] The use of Gauss-Hermite quadratures is convenient for taking expectations over log-normally distributed random variables, given that one needs only to scale, shift and exponentiate the original nodes.

especially beneficial since this routine is later nested within the dynamic estimation procedure. A 10-point one-dimensional quadrature is used.

This procedure iterates on values of $\alpha, \beta, \sigma_\varepsilon^2, \omega_0$ and σ_η^2 until the demand orthogonality conditions are met and the quantity moments are closest to being matched.

6.3 Firm capacity decisions

6.3.1 Formulation

Firms in the DRAM industry can change capacity levels mainly through two procedures. First, they resort to activation and deactivation of production lines in existing facilities. These decisions are used to address the foreseeable market conditions. Second, firms also resort to building and scrapping production facilities altogether. Such decisions often take more than two years to be carried out and depend on a number of factors: the long-term vision of the firm about the industry and the competition, the broader interests and the availability of financial resources to its parent company, its relative technology capabilities, the age of existing production facilities, etc. This paper focuses on the first type of capacity decisions. Such decisions are sensitive to the information about market conditions and are taken frequently. Because broader capacity decisions are influenced by a larger number of factors and are associated with long lag times, incorporating them would require a more extensive dataset and an additional level of analysis.

Firms take into account the present and future impacts of changing capacity levels. It is worth revisiting the Bellman equation for the firm's decision:

$$\begin{aligned} V_i\left(\Omega_t, \Theta_{it}\right) = \delta\, E\left[\pi^*_{it+1}\middle|\, \Omega_t, \Theta_{it}\right] \quad &+ \quad \underset{K_{it+2}}{max}\, \delta\, E_{\varepsilon_{t+2}, s_{-it+2}}\left[V_i\left(\Omega_{t+1}, \Theta_{it+1}\right)\middle|\, \Omega_t, \Theta_{it}\right] \\ &- \quad InvCost\left(K_{it+1}, K_{it+2}\right) \end{aligned} \tag{8}$$

Above, the equilibrium quantities have already been found through the procedure described in the previous section, and the Bellman equation can be used to understand the firms' capacity decisions. Given the short-term interpretation of the changes in capacity levels, the functional form of the cost of changing capacity levels is assumed to be

$$InvCost\left(K_{it+1}, K_{it+2}\right) = \omega_1 \phi^{t+1}\left(K_{it+2} - K_{it+1}\right)$$

where ω_1 and ϕ are parameters to be estimated and hold the following interpretation: ω_1 captures the constant unit cost of increasing capacity from level K_{it+1} to K_{it+2}, adjusted by ϕ, which captures factors that may change the cost of adding additional capacity over time. The linear specification for the cost of adding capacity was discussed with managers in the industry to be the most realistic one for modeling the pooled costs of changing short-run capacity levels. In addition, it simplifies the dynamics of the problem and make it estimable in a feasible amount of time. The idea is to transform the original problem into one that is easier to solve and estimate. In particular, the problem above can be recast into a another where the firm has to build up the full capacity level from scratch in each period, but is able to 'cash it' back after it has used it for production, at a lower value. To see this, consider the impact of the capacity decision K_{it+2} to the future: it influences the product competition stage in period $t+1$ (when quantities $t+2$ are decided) and it affects the cost of changing capacity levels from K_{it+2} to K_{it+3}. There exist no further implications of changing the level K_{it+2}. Because the payoffs are additively separable, when firm i chooses its capacity level it affects the expected firms' future payoffs by

$$\begin{aligned} \tfrac{\partial}{\partial K_{it+2}} \quad & E_{s_{-it+2}} \left(\max_{q_{it+2}} \delta^2 E_{\varepsilon_{t+2}} \left[\pi_{it+2} \middle| s_{it+2}, s_{-it+2}, c_{t+2} \right] \middle| s_{it+2} \right) - \omega_1 \phi^{t+1} \left(1 - \delta\phi\right) K_{it+2} \\ & s.t.\, q_{it+2} \leq K_{it+2} \end{aligned} \tag{9}$$

where the first term captures the expected equilibrium benefits of increasing the capacity constraint by one unit, while the second aggregates the investment costs related to changes in capacity levels, i.e., today's effect $-\omega_1\phi^{t+1}K_{it+2}$ and tomorrow's effect $+\delta\omega_1\phi^{t+2}K_{it+2}$. Problem (9) captures all dynamic effects of the capacity decisions K_{it+2}. This is because it does not depend on previous (or subsequent) capacity decisions. The linear specification is equivalent to the assumption that short-run capacity levels do not suffer from 'state-dependence' in the sense that their equilibrium levels are not related to the previous levels except through other state variables. Consequently, there is no need to consider additional implications of the capacity decisions because all payoff-relevant effects of changing K_{it+2} are already captured in formulation (9), which is the requirement of the Markov perfect equilibrium concept.[19] Figure 3 plots the

[19]Another way to see this is true is by noticing that under the linear cost specification the current capacity level drops from the first-order condition for the choice of future capacity.

equilibrium capacity policies recovered from solving the first-order conditions as a function of the private demand signals. They are approximated by piecewise linear functions.

Finally, notice that the incremental benefit of adding one unit to capacity is not equal to the shadow-cost of the firm's capacity constraint. Because the capacity level K_{it+2} is observable by all firms in the next period, firm i should expect strategic responses from its rivals in terms of future production levels. Recovering the shadow cost from quantity competition would provide information on the direct effect but would ignore further strategic effects.

6.3.2 Numerical Approach to the Equilibrium Solution

Because the individual information shocks s_{it+2} are unknown to the econometrician prior to the estimation of the problem, the solution strategy involves finding the functions K^*_{it+2} that solve the capacity first-order conditions (9) for all values s_{it+2}, for all firms at each time period. Hence, for each firm at each time period the solution to the capacity decision is given by a function $K^*_{it+2}(s_{it+2})$, $s_{it+2} \in \mathbb{R}_+$. Finding such equilibrium functions exactly would entail either having access to a closed-form solution or to be able to solve the set of capacity first-order conditions at an infinite number of values of the demand signals. Instead, the capacity policies are approximated by piecewise linear interpolation functions at D points, in the interval $[a, b]$. The piecewise linear function interpolation exhibited better stability properties than other candidate interpolation methods (such as C-splines and Akima splines). For each firm the optimal policy is evaluated at $D = 5$ demand signal points, in the domain [0.15, 3.0].[20] Those points are located at the Chebyshev nodes on that domain, so as to maximize the stability of the approximation to the policy functions.

The idea behind the numerical approach used to solve the asymmetric information equilibrium is to make sure the capacity first-order conditions are satisfied at each of the D domain points, for all firms. It entails solving a system of $nD = 20$ equations with respect to 20 interpolant control points in order to find the equilibrium capacity policies. Since the equilibrium expression of the expected profits is not analytical, finite differences are taken in order to

[20] A number of sensitivity tests were performed around these values with little impact. The interpretation of the bounds is that the equilibrium policies are searched over a domain region where the stochastic component can influence the market willingness-to-pay from roughly 15% to 300% of the deterministic component. Outside these bounds extrapolation would be needed. These are however generous bounds over the potential ignorance of the firm, and contain the amplitude of the demand shocks recovered in the static estimation.

calculate the numerical derivatives present in the capacity FOC's.

The goal of this stage is to recover policies $K^*_{it+2}(s_{it+2})$, $i \in n$, which are parametrized by vectors of 'y-values' at the Chebyshev nodes. For each firm, for each of the pre-chosen r^{th} nodes, $s^{(r)}_{it+2}$, a Gauss-Hermite quadrature $G_0\left(s_{-it+2}|\, s^{(r)}_{it+2}\right)$ is formed so as to take the outer expectation of (9). Hence, the strategy is to consider that if the firm received demand signal $s^{(r)}_{it+2}$, it is likely that the other firms received signals $G_0\left(s_{-it+2}|\, s^{(r)}_{it+2}\right)$, with the respective probabilities assigned by the quadrature. The firm will then solve the quantity equilibrium given that it decided capacity $K^*_{it+2}\left(s^{(r)}_{it+2}\right)$, and the other firms will choose capacities $K^*_{jt+2}(s_{jt+2})$, $j \neq i$, where s_{jt+2} is provided by $G_0\left(\cdot \left| s^{(r)}_{it+2}\right.\right)$. An intuitive way to think about this procedure is that firms are imagining which scenarios are likely to occur given their private information. The quadrature simply picks such 'optimal' scenarios that maximize the numerical efficiency of the expectation operator. After imagining each case, the firm will take expectations according to the weights associated with quadrature G_0. Hence, in order to solve the quantity stage, Kuhn-Tucker conditions are evaluated at the candidate capacity policies and the expectation $E_{\varepsilon_{t+2}}\left[\varepsilon_{t+2}|\sum_{i\in n} s_{it+2}\right]$ is formed through, where the sum is taken over the demand signals of each node of quadrature G_0 (3 quadrature nodes for the signals of the rival firms, plus the conditional firm i's node $s^{(r)}_{it+2}$) . Finally, quadrature G_0 is chosen to be a 3-dimensional 7^{th} degree quadrature, evaluated at $2^d + 2d^2 + 1 = 27$ nodes, so that precision is privileged while looking for profit increments as a result of small changes in capacity levels.

As before, parameters ω_1 and ϕ are estimated through matching the moments of the distribution of the sample capacity data with those of the model prediction. In particular the moments $E\left[K^l_{it} - \widehat{K}^{*l}_{it}\right]$, $l = 1..4$ are used, where $\widehat{K}^*_{it} = E\left[K^*_{it}(s_{it})|\, \widehat{\varepsilon}_t\right]$. In addition, the moment $E\left[(K_{it} - K_{it-1}) - \left(\widehat{K}^*_{it} - \widehat{K}^*_{it-1}\right)\right]$ is also introduced, so as to capture differences over time that might inform parameter ϕ. While ω_1 and ϕ are estimated separately from the static parameters, demand shocks that were recovered from the demand estimation are still used at this stage in order to increase the estimation efficiency. After the equilibrium policy functions for firm i at time t are recovered through solving the equilibrium, predicted capacities are generated by evaluating them at the correct signals centered at the recovered demand shocks. Finally, quadratures are also used to construct $\widehat{K}^*_{it}$, using the density $f_{s_t|\widehat{\varepsilon}_t}$.

7 Results

7.1 Static Parameters

The static parameters were estimated through an efficient 2-step generalized method of moments and the estimation results are presented in Table 4. The demand parameters reveal a price elasticity of 1.02. The parameter σ_{ε}^2 has an indirect interpretation since it measures the variance of the logarithm of the error term. The implied standard deviation of the lognormal distribution is of 31.1%. The interpretation for this number is that before using information to tackle the stochastic component of demand, firms face an uncertainty (measured by the standard deviation) of 31.1% of the systematic price. Table 5 presents the impact of the demand shock on price at different scenarios, described in terms of percentiles of the lognormal distribution. The "Price Impact" columns in Table 5 represent the percent change in price associated with the percentile of the probability of that change occurring. For example, with 10% probability market factors may decrease price to less than 64% (percentile 10%) or increase it to more than 40% of its mean level (percentile 90%).

The recovered market shocks are plotted in Figure 4. In line with the perception of the industry, the market showed poor performance near the end of 2008 and the beginning of 2009. The correlation between the market demand shocks and production and capacity levels is of 0.37 and 0.28 respectively. This establishes some relation between the firms' decisions and the market-level shocks. In addition, the fact that the correlation between market shocks with production is higher than that with capacity makes sense since production is a more flexible decision, while capacity is decided further in advance and managers have less information about market conditions when setting it. As an indication of the precision of the information estimated from the data, the posterior variance of the demand shocks when the short-run demand signals are equal to their means is[21]

$$V\left(\varepsilon_t \middle| s_t = 1\right) = e^{\frac{4\sigma_{\varepsilon}^2\sigma_{\eta}^2}{4\sigma_{\varepsilon}^2+\sigma_{\eta}^2}}\left(e^{\frac{\sigma_{\varepsilon}^2\sigma_{\eta}^2}{4\sigma_{\varepsilon}^2+\sigma_{\eta}^2}} - 1\right) \tag{10}$$

which is equal to 6.8%, i.e., the demand information reduces the price variance by 29.4%. Sim-

[21] The variance of the posterior log-normal depends on the actual values of the demand signals. Here, we evaluate the signals at their mean values (=1).

ilarly, the variance of the 'long-run' demand signals when the private market signal is equal to its mean is given by

$$V\left(\varepsilon_t | s_{it}=1\right) = e^{\frac{\sigma_\varepsilon^2 \sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2}} \left(e^{\frac{\sigma_\varepsilon^2 \sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2}} - 1\right) \tag{11}$$

which is equal to 8.77%, an uncertainty reduction (in terms of variance) of 9.1%.

Finally, the cost parameter reveals a wafer production cost of \$734.96. This number was discussed with managers in the industry and deemed in line with expectations.[22]

7.2 Dynamic Parameters

The parameter estimates associated with changes in capacity levels, of adding capacity, ω_1 and ϕ are presented in Table 6. As expected, increasing capacity is costly for the firm. In addition, it is cheaper to add the same amount of capacity later than earlier. The percent difference in the investment cost of adding the same amount of capacity tomorrow vs. today is given by:

$$\Delta\% InvCost = \frac{\omega_1 \phi^{t+2} K_{it+2} - \delta \omega_1 \phi^{t+1} K_{it+2}}{\delta \omega_1 \phi^{t+1} K_{it+2}} = \frac{\phi - \delta}{\delta}$$

which is equal to -15.13%, so that increasing capacity by 10 million gigabytes meant an investment of approximately \$4.88 billion in the second quarter of 2005, but only of \$100 million in the third quarter of 2010. This reflects the technological improvements achieved by the DRAM industry over time.

8 Counterfactual Analysis

Once the model fundamentals are estimated, one can consider alternative scenarios. The following section discusses the impact of the quality of the information on firms' (short-run) operating profits.

[22]The cost above is for a representative wafer, since the actual wafer cost also depends on its size, namely, 200mm or 300mm.

8.1 The Value of Market Information to Production Decisions

This section analyzes the effects of the information structure on the firms' profits, taking capacity decisions as given. The goal is to analyze the impact of the quality of market information on firm performance. I first describe the expected impact of the quality of information on the industry. Then, I analyze the outcomes for the market 'bust' period, from the fourth quarter of 2008 until the second quarter of 2009.

The impact of improving the precision of market information available to the industry on a firm's own profitability is ambiguous: if on one hand, the firm benefits from being able to make better decisions, on the other the fact that its rivals also have access to better information may put downward pressure on the firm's profitability. The size of the net impact depends on the actual parameters of the problem. For example, under the current assumptions symmetric firms - when unconstrained by capacity - benefit from information precision when facing an elastic demand curve. This effect disappears as the market elasticity becomes unitary.[23] Moreover, the aggregate effect of information can be separated into the cases when the market performs above and below average. Figure 5 decomposes the overall effect of the variance of information σ_η^2 on the industry's profits, in the cases of the underlying market demand shock being above or below average. As expected, the industry is more profitable when the market performs better. More interestingly, market information is beneficial for firms when the market performs below average, but is damaging whenever the market performs above average. Because the elasticity of the demand estimated from the data is close to one, the overall effect shown in Figure 5 is only slightly positive. These effects change considerably once capacity constraints are introduced.

The decomposition shown in Figure 6 reveals that information can have a higher impact on industry profits once capacity constraints are taken into account. This happens because capacity limits production exactly when the market does better, and so the negative role of the precision of information is attenuated. The average effect of the quality of information is now visibly higher, as shown in Figure 6. Capacity levels may also attenuate the role of information: when they are low enough, producers are very likely to produce up to capacity independently of the information, and so the quality of information may matter less once again.

What is the value of market information for short-run decisions in the DRAM industry?

[23] A proof is available in the appendix.

The answer to this question depends on a number of variables, such as the installed capacity, the production costs, etc.. Figure 7 plots the ex-ante expected profits from the first quarter of 2008 to the third quarter of 2010 for the industry under three conditions: perfect information, estimated quality of information and no information. Respectively, these scenarios correspond to $\sigma_\eta^2 = 0$, $\sigma_\eta^2 = \widehat{\sigma_\eta^2}$ and $\sigma_\eta^2 = \infty$.[24] The profit levels shown are ex-ante in the sense that the market shock is not known *a priori*. First, notice that expected profits have been lowering over time. This is attributable mainly to capacity levels: the higher the installed capacity, the more likely firms are to reach the unconstrained Cournot equilibrium, which yields less profits than the case where production is limited. This is a consequence of the Cournot equilibrium exhibiting a *prisoner's dilemma*, where firms would be better off by jointly decreasing production levels, but do not have unilateral incentives to do so. Hence, low capacity levels have the potential to solve the dilemma. Second, the expected industry profits are generally higher when firms have perfect information available, and lower when no information is available. Table 7 decomposes the ex-ante value of information for the selected periods: First, consider the last set of columns (columns 6 and 7), which measure the potential value of information in the market (additional profits from absence of marketing information to perfect information). The potential value of information in this period goes from approximately \$17 million in the first quarter of 2008 to approximately \$369 million in the second quarter of 2009. Of these, firms expected to appropriate \$1.59 million in the first quarter of 2008, and approximately \$203 million in the second quarter of 2009, for example. Finally, during 2009 and the first three quarters of 2010, the industry expected to increase its profits by roughly 7.39% out of a potential increase of 15.48%, due to the use of market information.

It is also interesting to analyze the market period crisis *ex post* (once demand shocks are known to the econometrician), from the fourth quarter of 2008 until the second quarter of 2009. Figure 8 shows that the actual operating profits were between the cases of perfect information and absence of information, conditional on the realized demand shocks. During that period the industry profits were \$1,87 billion above the case of absence of information, and would have increased an additional \$1,91 billion if perfect information were available. Finally, the existence

[24]Figure 7 starts in the first quarter 2008. Before this quarter, profits were generally higher and there are little differences in scenarios across conditions. This is due to the mechanism that was just exposed, of how very low levels of capacity rendering the precision of information to have little effect. Also, the scenario for $\sigma_\eta^2 = \infty$ was simulated with $\sigma_\eta^2 = 10.0$ which was found to be high enough so that information has negligible impact.

of market information prevented firms from producing below variable cost in the last quarter of 2008 and the first quarter of 2009.

8.2 The Value of Market Information

Unlike production decisions, capacity decisions are made under private, correlated information. It follows that the effects of information on firms' profits is not trivial. The available literature provides a rule of thumb to assess the effects of information. It relates the correlation induced by having better information with the slope of the best-response curves of the control variable. The current analysis is complex for two reasons: first, we have no direct theoretical guide or access to an analytical solution to understand the effects of the quality of information on capacity decisions. Second, although the intuition about the role of the slopes of the best-response curves of production could carry over to capacity decisions, they cannot be used directly since those best-response curves are non-linear when firms face a constant-elastic demand (see Figure 9). Hence, the findings presented below are applicable to the current industry under the estimated parameters, but cannot be generalized to other contexts and values of estimates.

Figure 10 summarizes the effect of the precision of information on expected profits for the industry. First, more precise information is always better than no information. In addition, perfect information is always better than the estimated precision of information only in the second quarter of 2010. For the remaining periods (including the ones before 2008) I find a non-linear effect of information precision, where the quality of information has a non-linear effect on firm profits. Table 8 summarizes the effects of the quality of information, including the effects of investment in capacities. The information available increases profits by up to 15.6% in the last quarter of the sample, and as little as 2.2% in the quarter before the last. The gains of information found here are higher than the ones comprising only the operating profits in Table 6. In addition, the total value of the *expected* gains from the existing information for the whole sample (i.e., since the second quarter of 2005, using the same discount factor of 0.95 yearly) is of $4,76 billion.

Finally, Figure 11 shows that firms would have benefited from better information during the crisis period, after incorporating the payoffs related to investment/divestment in capacity. In effect, better information would have created a divestment such that firms would have been better off during the crisis period by reducing capacities, and effectively gaining and additional \$4,92 billion. Still, the value of the existing market information is estimated to have been worth \$1,65 billion during this period. Finally, the value of the available information for the whole sample (excluding the first quarter) conditional on the realized market shocks yielded returns of \$6,38 billion, or approximately \$290 million for the industry per quarter.

9 Conclusion

Demand information is a fundamental component for decision-making in competition settings. This is especially true in industries where decisions are made in advance of actual demand levels becoming known. This paper takes the example of the DRAM industry. It presents a model that allows firms to setup capacity and production levels, and estimates the precision of the market information available to managers and the value of information in terms of industry profits. The short-run information is found to have a positive impact on the returns from quantity competition. However, a non-linear effect on profits is found when capacity decisions are taken into account, and an (interior) 'optimal' level for the quality of information is uncovered.

Market information is found to play a role in the performance of the industry: it is responsible for \$1,65 billion in operating profits during the recent DRAM 'bust' from the fourth quarter of 2008 until the second quarter of 2009. In addition, the existing market information was valued at \$6,38 billion for the whole sample, an average of approximately \$290 million for each quarter.

Market information has potential to generate further gains for the industry. Being able to anticipate downturns has as an especially high impact on the industry performance: during the late 2008 market crisis firms could have earned an additional \$4,92 billion if perfect information had been available. This was a combination of different effects: the fact that the industry benefits from precise information especially when the market performs badly, and that at the time capacity was high enough so that anticipating necessary divestment would have yielded benefits for the firms.

Finally, some future avenues for research include relaxing the parametric assumptions associated with the information structure and incorporating heterogeneity in the quality of information available to different firms.

10 Figures

Figure 2: Equilibrium Quantities as a function of the Expected Market Shock

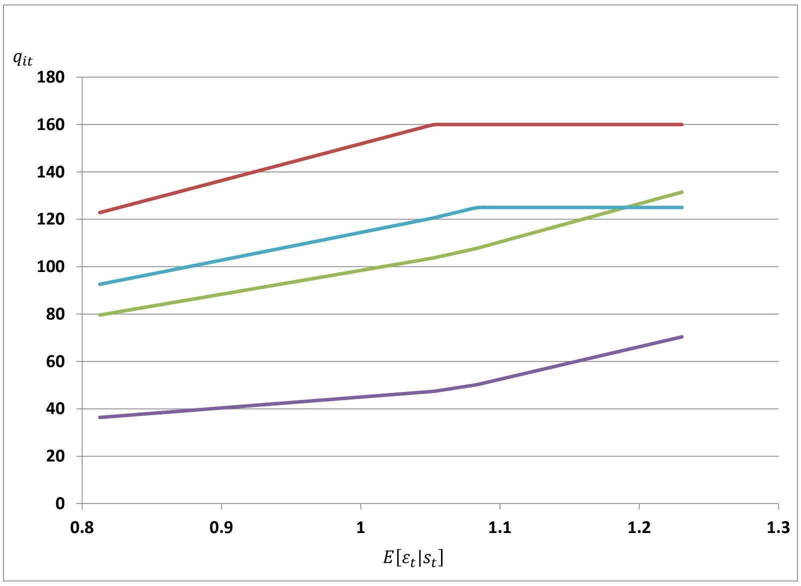

Production levels are increasing on the expected market performance up to installed capacity levels. This plot uses the estimated parameters, as well as the cost data for a particular time period.

Figure 3: Equilibrium Capacity Policies as a function of the Private Demand Signal

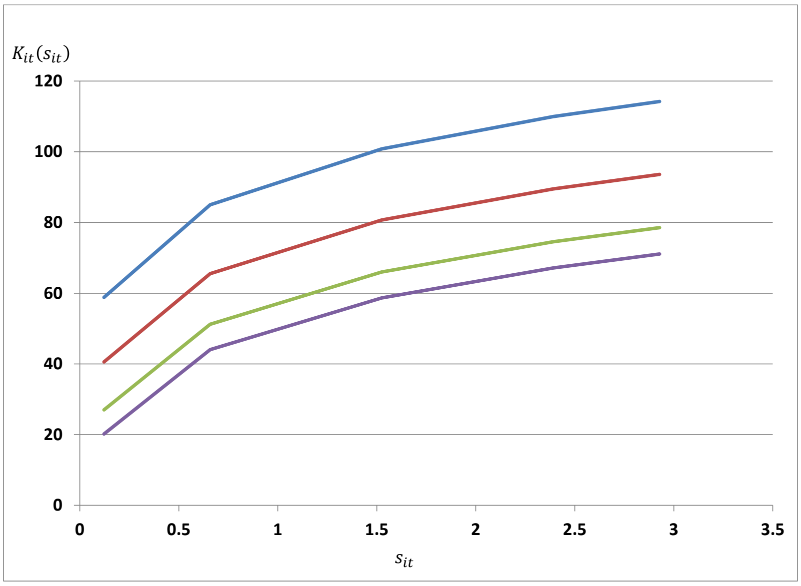

The equilibrium capacity levels are increasing on private demand signals and are approximated by piecewise linear functions. This plot uses the estimated parameters as well as the cost data for a particular time period.

Figure 4: Recovered Demand Shocks (%)

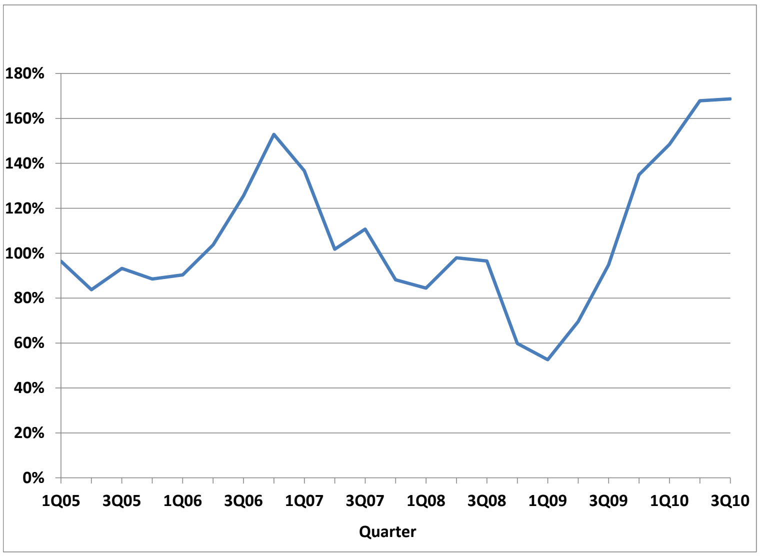

As expected, the recovered demand shocks are centered at 100%, having the most positive impact on willingness-to-pay in late 2006 and late 2010, and affecting the market most negatively near the end of 2008 and beginning of 2009.

Figure 5: Expected Industry Profits Conditional on Market Shocks

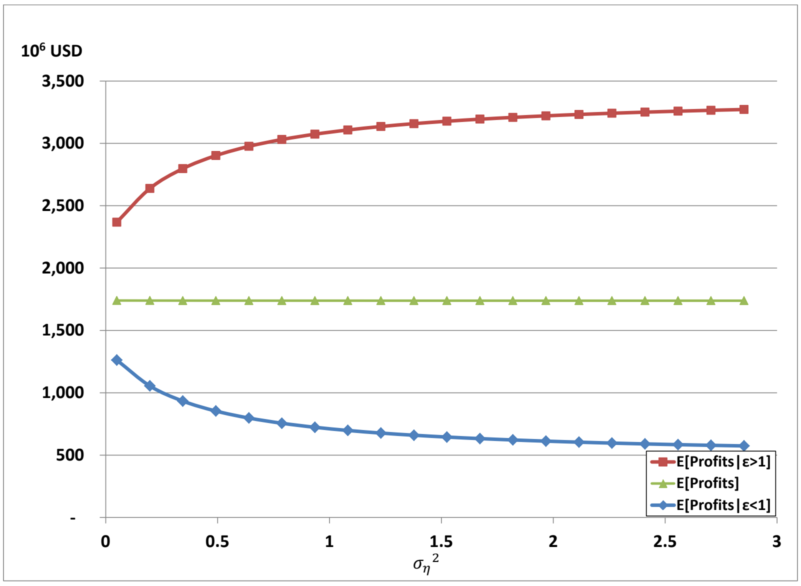

More precise market information ($\sigma_\eta^2 \to 0$) is better whenever the underlying market shock ε is lower than one, but worse when the market shock is higher than one. This plot uses the estimated parameters, as well as the cost data for a particular time period.

Figure 6: Expected Industry Profits with Capacity Constraints, Conditional on Market Shocks

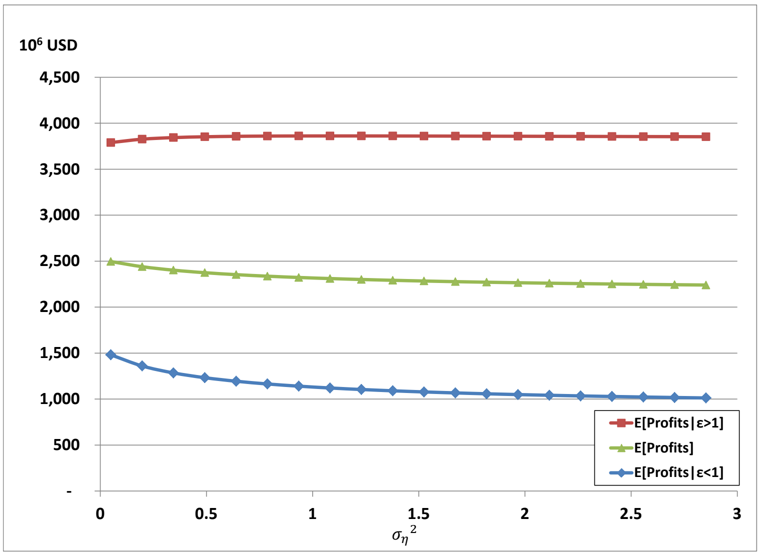

The impact of the precision of information on the expected profits when the market performs above average is now flatter than before. This is due to capacity constraints that limit production levels from above, and soften the role of information when the market performs better.

Figure 7: Expected Industry Operating Profits for 1Q 2008 - 3Q 2010

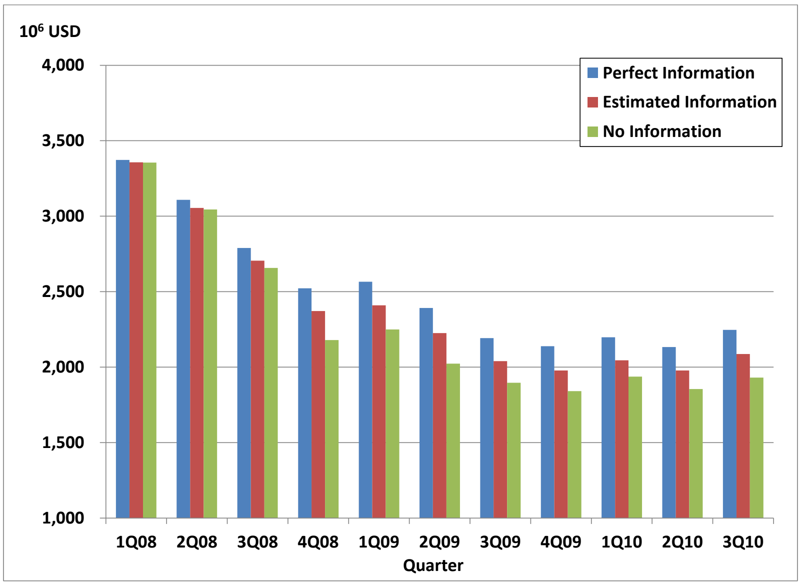

The expected industry profits given the estimated precision of market information is between the perfect information case and the case of absence of information. The quarters before 2008 showed little difference in profitability depending on the information scenarios, due to the existence of very low capacity levels.

Figure 8: Industry Operating Profits Conditional on Market Shocks for 4Q 2008 - 2Q 2009

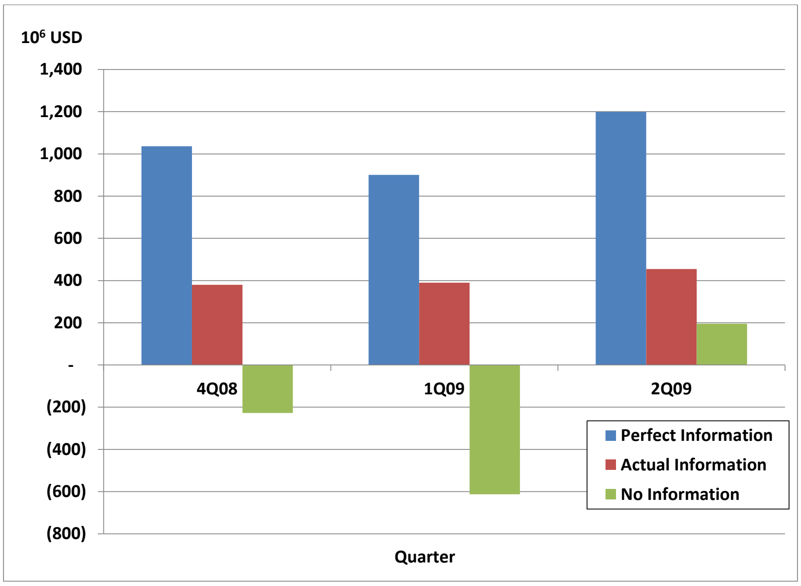

During the market crisis period, the total (expected) benefit for firms from taking market information into account was of $1,87 billion. The scenario of perfect market information would have yielded an additional $1,91 billion.

Figure 9: Best-Response Curves for Cournot Duopoly, under no Capacity Constraints

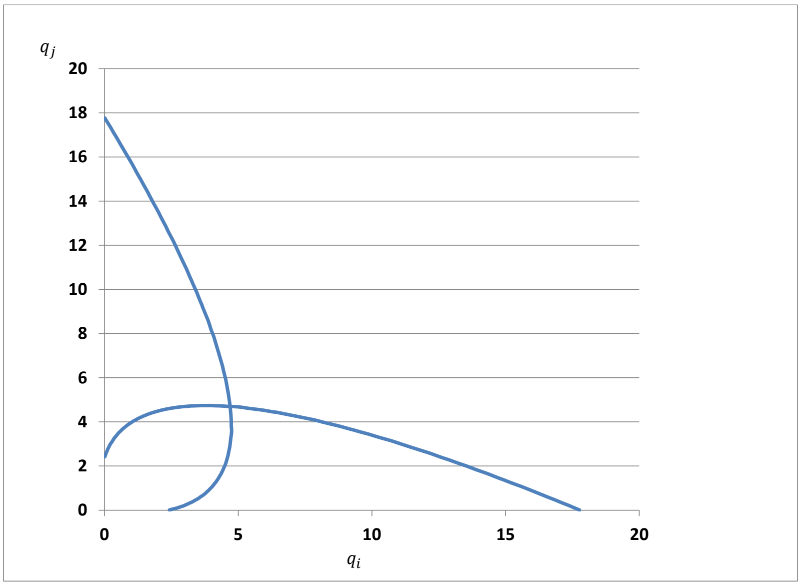

When firms face a constant-elastic demand curve, best-response curves become non-linear. In particular they are positively-sloped for small quantities, and negatively-sloped for higher quantities. Values: $\alpha = 100$, $\beta = -0.8$, $c_1 = c_2 = 10$.

Figure 10: Expected Industry Operating Profits for 1Q 2008 - 3Q 2010, including Investments in Capacities

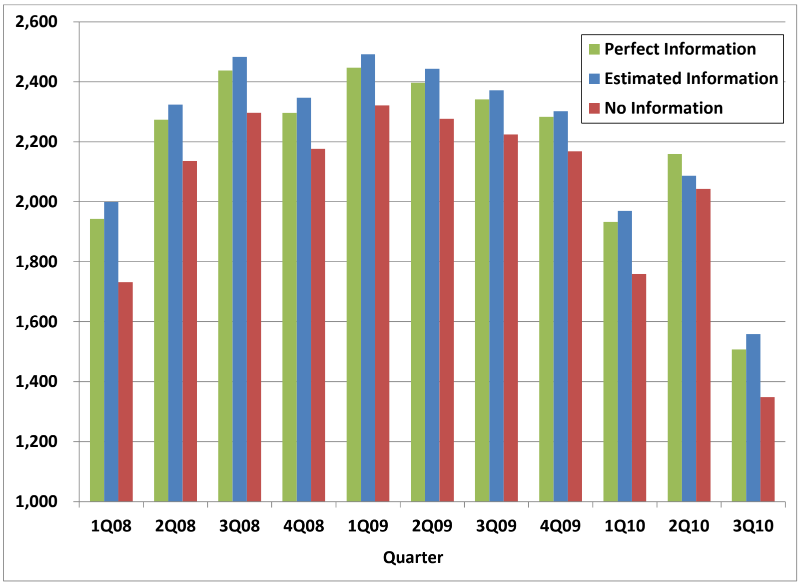

More precise information is always better than its absence. In addition, the quality of information often has a nonlinear effect on profits, where an intermediate quality of information is optimal.

Figure 11: Industry Operating Profits Conditional on Market Shocks for 4Q 2008 - 2Q 2009, including Investments in Capacities

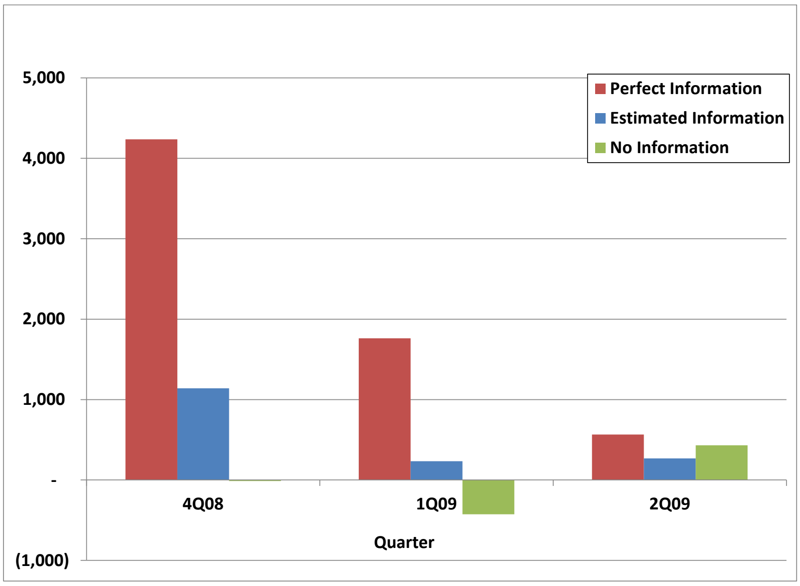

During the market crisis period, the total (expected) benefit for firms from taking account market information was of $1,65 billion. The scenario of perfect market information would have yielded an additional $4,92 billion. The high numbers presented above incorporate the returns from divesting in capacity given the poor performance of the market.

11 Tables

Table 1: Shift in Geographical Distribution of DRAM Production

Region	Number of DRAM Makers		Market Shares	
	1992	2009	1992*	2009
US	3	1	5-8%	11.0%
Japan	12	1	65-70%	2.4%
Korea	3	2	22-25%	50.6%
Europe	2	1	2%	9.6%
Taiwan	–	3	–	8.1%

Distribution of DRAM producers by geography, and world unit market shares by the end of 1992 and in February 2009 (only main firms, amounting to 95% market share). * - 1992 shares refer to the 1Mb chip market. Source: Murillo (1993), and Gartner Inc. as cited at fabtech.com.

Table 2: Descriptive Statistics of the Data

	Min	Max	Mean	Std. Dev.
Avg. Yearly Industry Sales (10^6 GB)	42.16	480.86	201.76	139.43
Market Price (\$/GB)	12.07	132.36	52.98	39.04
Avg. Yearly Industry Capacity (10^6 GB)	42.16	480.86	210.99	146.71
Efficiency Indexes	0.69	2.00	1.28	0.38

Observations: 23 Periods, 4 Alliances
Efficiency indexes are relative to a normalized value for alliance 1 at a specific period in time.

Table 3: DRAM Firms in Sample, Nationalities and Alliances

Name	Nationality	Alliance Identifier
Samsung	South Korea	1
Hynix	South Korea	2
Powerchip	Taiwan	3
ProMOS	Taiwan	3
Elpida	Japan	3,4
Micron	USA	4

Table 4: DRAM Demand Parameter Estimates

Parameter	Coefficient	Standard Error
$\alpha^{'}$	8.580	0.445
β	0.978	0.096
σ_{ε}^2	0.092	0.001*
ω_0	0.073	0.010
σ_{η}^2	0.510	0.167

* - The standard error of σ_{ε}^2 was calculated based on the analogous MLE linear estimation.

Table 5: Demand Shock Percentiles and Impact on Market Price

Percentile	Price Impact	Percentile	Price Impact
10%	64.73%	60%	103.13%
25%	77.83%	75%	117.19%
40%	88.44%	90%	140.89%

Note: the size of the asymmetries of positive and negative shocks above is a combination of the lognormal assumption and of the estimated variance of the log-error term, $\widehat{\sigma}_{\varepsilon}^2$.

Table 6: Dynamic Parameter Estimates

Parameter	Coefficient	Standard Error
ω_0'	0.520	0.131
ϕ	0.823	0.278

Above, $\omega_0' \equiv 1000.\omega_0$ so that the order of magnitude of the parameters estimated is the same.

Table 7: The Value of Short-Run Information ($\$10^6$)

Quarter	Perfect Inf. − Est. Inf.	% Diff.	Est. Inf. − No Inf.	% Diff.	Perfect Inf − No Inf	% Diff.
1Q08	15.63	0.47%	1.59	0.05%	17.22	0.51%
2Q08	53.20	1.74%	9.88	0.32%	63.08	2.07%
3Q08	84.20	3.11%	47.39	1.78%	131.59	4.95%
4Q08	150.32	6.34%	191.76	8.80%	342.08	15.69%
1Q09	155.90	6.47%	159.92	7.11%	315.82	14.04%
2Q09	166.08	7.46%	202.87	10.03%	368.95	18.24%
3Q09	153.31	7.52%	143.01	7.54%	296.32	15.63%
4Q09	160.88	8.13%	136.85	7.43%	297.73	16.17%
1Q10	152.82	7.47%	107.58	5.55%	260.40	13.44%
2Q10	154.56	7.81%	123.30	6.65%	277.86	14.98%
3Q10	159.28	7.63%	156.87	8.13%	316.15	16.38%

Each column displays the difference in industry profits between two scenarios. Note that the third set of columns captures the aggregate effect of the first two sets.

Table 8: The Value of Information($\$10^6$)

Quarter	Perfect Inf. − Est. Inf.	% Diff.	Est. Inf. − No Inf.	% Diff.	Perfect Inf − No Inf	% Diff.
1Q08	(55.30)	-2.77%	267.36	15.44%	212.06	12.25%
2Q08	(50.67)	-2.18%	188.66	8.83%	137.99	6.46%
3Q08	(45.17)	-1.82%	186.66	8.13%	141.49	6.16%
4Q08	(50.38)	-2.15%	170.13	7.82%	119.75	5.50%
1Q09	(44.81)	-1.80%	170.54	7.35%	125.73	5.42%
2Q09	(46.56)	-1.91%	166.80	7.33%	120.24	5.28%
3Q09	(30.44)	-1.28%	147.41	6.63%	116.97	5.26%
4Q09	(18.82)	-0.82%	133.63	6.16%	114.81	5.29%
1Q10	(36.86)	-1.87%	210.79	11.98%	173.93	9.89%
2Q10	71.64	3.43%	44.16	2.16%	115.80	5.67%
3Q10	(50.96)	-3.27%	209.70	15.55%	158.74	11.77%

Each column displays the difference in industry profits between two scenarios. Note that the third set of columns captures the aggregate effect of the first two sets.

12 Appendix

Figure A: 12-inch (300 mm) Silicon Wafer

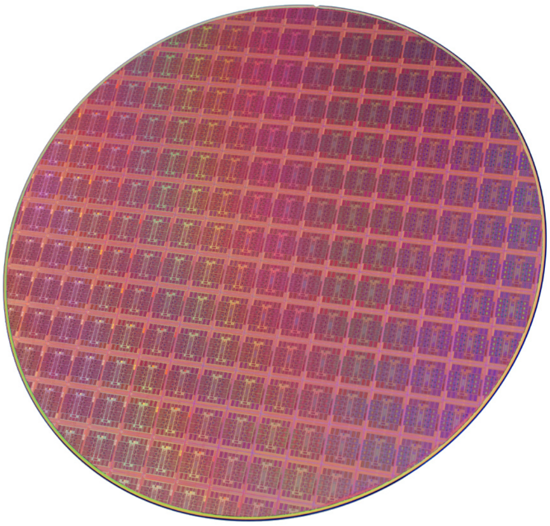

Silicon wafers are 'blank slates' for several semiconductor industries, like RAM, Flash memory, CPU's, etc.

12.1 Conditional Distribution

In order to determine the posterior distributions of one demand signal given the demand shock and remaining demand signals, one can use the equality

$$log(s_{it}) = log(\varepsilon_t) + log(\eta_{it}).$$

Since the logarithm of a lognormally distributed random variable follows a normal distribution, one can setup the conditional distributions by use of the equality above and the explicit conditional and posterior distributions. The posterior distribution of the demand shock and the competitors' demand signals given firm 1's demand signal (without loss of generality) becomes:

$$
\begin{array}{c} \varepsilon \\ s_{2t} \\ \vdots \\ s_{Nt} \end{array} \Bigg|_{s_{1t}} \sim LogN \left(\begin{bmatrix} log(s_{1t})\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2+\sigma_\eta^2} \\ log(s_{1t})\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2+\sigma_\eta^2}-\frac{\sigma_\eta^2}{2} \\ \vdots \\ log(s_{1t})\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2+\sigma_\eta^2}-\frac{\sigma_\eta^2}{2} \end{bmatrix}, \begin{bmatrix} \frac{\sigma_\varepsilon^2\sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2} & \frac{\sigma_\varepsilon^2\sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2} & \cdots & \frac{\sigma_\varepsilon^2\sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2} \\ \frac{\sigma_\varepsilon^2\sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2} & \sigma_\eta^2+\frac{\sigma_\varepsilon^2\sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2} & \cdots & \frac{\sigma_\varepsilon^2\sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\sigma_\varepsilon^2\sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2} & \frac{\sigma_\varepsilon^2\sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2} & \cdots & \sigma_\eta^2+\frac{\sigma_\varepsilon^2\sigma_\eta^2}{\sigma_\varepsilon^2+\sigma_\eta^2} \end{bmatrix} \right) \tag{12}
$$

which suggests that the knowledge of one's own demand signal is informative about the demand shock as well as about the signals of the competitors..

12.2 Effect of Market Information on Firm Profits' when $\beta = 1$

Here I prove (for the symmetric case of four firms, but generalizations are straightforward) that the quality of information is irrelevant under unitary demand.

Let $E_{\varepsilon|s}[\pi_i] = \max_{q_i} E_{\varepsilon|s}\left(\alpha Q^\beta \varepsilon - c \middle| s_i\right) q_i$, where $\varepsilon \sim LogN\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon\right)$ and $s_i = \varepsilon.\eta_i$ and $\eta_i \sim LogN\left(-\frac{\sigma_\eta^2}{2}, \sigma_\eta\right) \forall i \in n$. Finally, s is a collection of signals $s_i, i = 1..n$. The first-order condition for firm i becomes:

$$
\alpha Q^{\beta-1}(Q+\beta q_i)\exp\left\{\frac{\sigma_\varepsilon^2\left(\sum_i \log(s_i)+2\sigma_\eta\right)}{4\sigma_\varepsilon^2+\sigma_\eta^2}\right\} - c = 0
$$

I inspect the symmetric case where $q_i = q_j$, $i, j \in n$. Solving for q_i and plugging it back into the expected profit function for firm i yields the expected profit conditional on the signal s_i. Taking expectations w.r.t s_i, i.e., $E_s\left[E_{\varepsilon|s_i}[\pi_i]\right]$ is done by noticing that $z \equiv \sum_i \log(s_i) \sim N\left(-2\left(\sigma_\varepsilon^2+\sigma_\eta^2\right), \sqrt{16\sigma_\varepsilon^2+4\sigma_\eta^2}\right)$ and solving:

$$
E_s\left[E_{\varepsilon|s}[\pi_i]\right] = \int_{-\infty}^{\infty} E_{\varepsilon|z}[\pi_i] f_z(z)\, dz = -\beta\alpha^{-\frac{1}{\beta}} 4^{\frac{1-\beta}{\beta}} c^{\frac{1+\beta}{\beta}} (4+\beta)^{-\frac{1+\beta}{\beta}} \exp\left\{\frac{2(1+\beta)\sigma_\varepsilon^2}{\beta^2\left(4\sigma_\varepsilon^2+\sigma_\eta^2\right)}\right\}. \tag{13}
$$

Finally, the gains of information are given by taking the comparative statics of (13) with respect to σ_η^2, which reveal that information precision is only strictly beneficial when $-1 < \beta < 0$, is irrelevant whenever $\beta = -1$ and is harmful whenever $\beta < -1$. Hence, profits only increase with information precision when demand is elastic, which is the case estimated from the data.

References

Aguirregabiria, V., and P. Mira (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75(1), 1–53.

Armantier, O., and O. Richard (2003): "Exchanges of Cost Information in the Airline Industry," *The RAND Journal of Economics*, 34(3), 461–477.

Bajari, P., C. L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75(5), 1331–1370.

Bajari, P., V. Chernozhukov, D. Nekipelov, and H. Hong (2009): "Semiparametric Estimation of Dynamic Games of Incomplete Information," *Working Paper*.

Chen, Y., C. Narasimhan, and Z. J. Zhang (2001): "Individual Marketing with Imperfect Targetability," *Marketing Science*, 20(1), 23–41.

Chu, W. (1992): "Demand Signalling and Screening in Channels of Distribution," *Marketing Science*, 11(4), 327–347.

Clarke, R. N. (1983): "Collusion and the Incentives for Information Sharing," *The Bell Journal of Economics*, 14(2), 383–394.

Fershtman, C., and A. Pakes (2009): "Finite State Dynamic Games with Asymmetric Information: A Framework for Applied Work," *Working Paper*.

Gal-Or, E. (1985): "Information Sharing in Oligopoly," *Econometrica*, 53(2), 329–343.

——— (1986): "Information Transmission–Cournot and Bertrand Equilibria," *The Review of Economic Studies*, 53(1), 85–92.

ICE-Corporation (1997): *Cost effective IC manufacturing 1998 to 1999*. Scottsdale, AZ: Integrated Circuit Engineering Corporation.

Jin, J. Y. (2000): "A Comment on 'A General Model of Information Sharing in Oligopoly,' Vol. 71 (1996), 260-288," *Journal of Economic Theory*, 93(1), 144 – 145.

Lariviere, M. A., and V. Padmanabhan (1997): "Slotting Allowances and New Product Introductions," *Marketing Science*, 16(2), 112–128.

Maskin, E., and J. Tirole (2001): "Markov Perfect Equilibrium: I. Observable Actions," *Journal of Economic Theory*, 100(2), 191 – 219.

Murillo, L. E. (1993): "The International Dynamic Random Access Memory Industry from 1970 to 1993 examined under the Dynamic Capabilities Prism: implications for technology policy," Ph.D. thesis, University of California, Berkeley.

Norton, J. A., and F. M. Bass (1987): "A Diffusion Theory Model of Adoption and Substitution for Successive Generations of High-Technology Products," *Management Science*, 33(9), pp. 1069–1086.

Novshek, W., and H. Sonnenschein (1982): "Fulfilled Expectations Cournot Duopoly with Information Acquisition and Release," *The Bell Journal of Economics*, 13(1), 214–218.

Pesendorfer, M., and P. Schmidt-Dengler (2003): "Identification and Estimation of Dynamic Games," *Mimeo, London School of Economics.*

Raith, M. (1996): "A General Model of Information Sharing in Oligopoly," *Journal of Economic Theory*, 71(1), 260 – 288.

Ryan, S. (2009): "The Costs of Environmental Regulation in a Concentrated Industry," *Working Paper.*

Schelling, T. C. (1960): *The strategy of conflict.* London: Oxford University Press.

Shapiro, C. (1986): "Exchange of Cost Information in Oligopoly," *The Review of Economic Studies*, 53(3), 433–446.

Villas-Boas, J. M. (1994): "Sleeping with the Enemy: Should Competitors Share the Same Advertising Agency?," *Marketing Science*, 13(2), 190–202.

Vives, X. (1984): "Duopoly information equilibrium: Cournot and bertrand," *Journal of Economic Theory*, 34(1), 71 – 94.