Title
The Low-Profit Trap in Hard Disk Drives, and How to Get Out of It

Permalink
https://escholarship.org/uc/item/1qw154f9

Author
Bohn, Roger E.

Publication Date
2000-04-01

marketplace

The Low-Profit Trap in Hard Disk Drives, and How to Get Out of It

Roger E. Bohn, Director, Information Storage Industry Center, UCSD

It is no news that the hard disk drive (HDD) industry suffers from chronically low profits. In bad years, most participants lose money, and even in boom years profits are not high. This situation has persisted for more than a decade, even as there is turnover among the second-tier participants. It's quite unusual for an industry to be in this situation over such a long period.

This article discusses the reasons for this chronic low profitability, and some modest suggestions on how to escape. The problem is a structural one at the industry level, not the individual firm level. Individual firms can lower their costs and improve the attractiveness of their products. But such improvements are quickly duplicated, so that they do not translate into sustained profits. In the next section, I sketch the industry conditions that lead to chronically low profits. I show what is needed in principle to escape. I speculate on two specific initiatives that would be steps in the right direction.

Why low profits, year after year?

Price and profits are determined in the short run by supply and demand. But profits in the long run are determined mainly by industry structure and characteristics. The hard disk industry has the structure of a commodity industry. In particular, the desktop HDD market:

- Products are undifferentiated, at least in the eyes of most purchasers.

- Firms compete mainly on price and dollars per gigabyte, not on features or other product attributes.

- Buyers have low-switching cost among HDD suppliers, and no customer loyalty or brand sensitivity.

- Entry barriers are low. A number of third-tier or prospective firms can rapidly build significant capacity, using open-market technology and hiring from incumbents.

CHASE H&Q   Hambrecht and Quist LLC

Statistical Summary Selected Data Storage Stocks

Industry Sector	Company	Shares Out. (MM)	Stock Price 2/1/00	Calendar Year Earnings Per Share			Price/Earnings Ratio		Market Cap ($MM)	Cal 99 Revs. ($MM)	Market Cap/ 99 Revs
				1998A	1999A	2000E	1999	2000			
Disk Drives	Maxtor Corp.	103.3	7.41	0.52	(0.82)	(0.53)	NM	NM	765	2,615	0.3
	Quantum HDD Group	86.0	7.94	(1.13)	(1.06)	0.30	NM	26.5x	683	3,415	0.2
	Seagate Technology	223.1	41.00	0.62	1.00	0.70	41.0x	58.6x	9,147	6,700	1.4
	Western Digital	121.1	4.69	(4.28)	(3.46)	(1.20)	NM	NM	568	2,072	0.3
Components	HMT Technology	45.7	3.22	0.37	(1.08)	(0.62)	NM	NM	147	204	0.7
	Hutchinson Tech.	30.7	16.63	(1.29)	0.02	(0.10)	NM	NM	511	530	1.0
	Komag, Inc.	65.4	2.53	(3.51)	(1.95)	(0.58)	NM	NM	165	273	0.6
	Read-Rite Corp.	49.8	2.78	(2.75)	(4.11)	(1.55)	NM	NM	138	640	0.2
Capital Equipment	Intevac	11.7	6.00	0.00	(0.98)	(0.60)	NM	NM	70	60	1.2
	Veeco Instruments **	17.5	58.75	1.08	1.45	1.84	40.5x	31.9x	1,028	300	3.4
Removable	Iomega Corp.	269.7	4.13	(0.15)	0.05	0.24	82.6x	17.2x	1,114	1,500	0.7
						Average:	41x	45x			0.8

**No official H&Q coverage; First call estimates. "A"= actual; "E"= estimate

- *Although there is important proprietary technology, it diffuses to other firms rapidly.*

- *There are fixed costs for R&D, but little economy of scale in manufacturing, so that being bigger does not give a major cost advantage.*

In second-year economics we learn that under these conditions, firms will enter the market until prices are driven down to "zero economic rents in the long run," meaning that accounting profits will only be enough to cover the cost of capital. We say that "new entrants compete away any economic rents," and that the number of competitors in the market will be large enough to approach the conditions of perfect competition.

Economic theory tells us that the only sustained way to get out of commodity-industry situations is to have product differentiation and barriers to entry. Product differentiation allows firms to compete on more than price, and divides the market into many segments with fewer competitors in each. Barriers to entry keep other firms out of each segment, at least for a few years. The stronger the product differentiation, and the stronger the barriers to entry, the more profitable the industry will be.

Escape?

Is it possible to escape from this situation? I think so, but it will take some radical changes. The dominant paradigm for progress in the industry has to change from lower cost per megabyte, via higher density per square inch. This has been, of course, spectacularly successful at the technical level. Density progress has at least kept up with, and for the last five years surpassed, the vaunted (and much misquoted) Moore's Law of semiconductors. Recently, with the potentially explosive growth of the non-PC and plummeting prices of PCs, there has been an additional imperative of lower cost per drive, which has also succeeded to an extent.

This paradigm has been so successful that it is now obsolete. Storage capacity has gotten ahead of demand. If the industry is to move to sustained profitability, it must move to new paradigms. I propose that cost per megabyte, and even cost per spindle, are no longer critical in many markets. Now the emphasis should be on perceived and achieved end-user value per dollar spent. HDDs should be designed and marketed as something that provides value to end users in their lives; not as an anonymous box that hides somewhere in their machine, and is noticed only when it causes problems. The goal is to create more product differ-entiation among mass storage systems, in the eyes of end users (and corporate purchasers). This will allow firms to mark out distinct market segments that they can dominate, and command a few percent more in price for. This would be enough to move the firms on the edge, to solid profitability.

I will discuss two examples of ways to increase end-user perceived value. The first is "desktop RAID," which is a specific case of a more general approach: moving downstream in the value chain. The second is to learn from Intel, which has managed to take a trivial few percent delta in PC performance, and turn it into something that end users will pay hundreds of dollars for. The HDD industry can do this with other improvements, some substantive and others only perceived.

What about HDDs in non-computer applications, such as digital video recorders, and HDDs in digital cameras? Those are useful long-term ways of increasing demand, and I think they are important for several reasons. But they will increase profitability in the long run only if some firms stay out of these segments. Since many others have written these new markets, I will say no more about them here.

Desktop RAID and Server in a Brick

HDD prices have now reached 1 cent per megabyte and $10 per gigabyte. For example, in October 1999, Fry's advertised an anonymous 13GB drive for $120. An IBM 20GB drive was about $190. At these prices for hardware, the cost of the HDD itself, becomes minor, except in entry-level systems. The human cost of gathering, validating, inputting, and just plain locating the data are many times larger than the cost of the raw HDD hardware. In essence, for many users, raw gigabytes of spinning drive are virtually free relative to other costs.

This has many consequences. For example, I expect a rapid growth in RAID mirroring (i.e. two distinct drives containing identical data). For anyone who has had a drive failure, it's worth $150 to have a continuous backup. Many of these failures are not actually due to the HDD, but in those cases a hot spare would prove useful, if only for debugging. Clever operating system (OS) design will make spares still more useful. For example, the second drive might be lagged behind the first one, in such a way that it provides "instant undo" if the OS or an application gets trashed.

A closely related concept is a "server in a brick," for the SOHO (small-office, home-office) market. Take a 10 to 20GB HDD; replace its EIDE interface with a 100Mbs ethernet; augment or substitute the microcontroller with a $10 RISC chip running Linux; add a single-chip 16MB RAM, a power supply, and a case, and voila—a server in a brick. Throw in $50 for service, support, and profit, and the box should be less than $250 at wholesale. At that price, a SOHO would buy two or three, hook them into the office ethernet, and run them mirrored—if the user interface were friendly enough. If one fails, you unplug it, take it back to the store for a $50 trade-in credit, and buy a new one. I won't repeat the arguments about whether thin PCs makes sense, but with this configuration the PCs could certainly get away with a small hard drive. If these devices were marketed as the Zip drive was, vendors could establish a brand name and consequent sustained profits.

Jim Gray of Microsoft has gone much farther in his projections. He envisions progress in both magnetics and silicon continuing, so

that in five years the normal configuration of an HDD is 128MB of memory, 100GB, a 200MHz processor, and it's effectively an application server in its own right. For example, drives could be sold with 20GB of data and applications on the drive, in encrypted form.

The theme here is to substitute cheap hardware and software for expensive end-users' and system-administrators' time. In the past the HDD companies ceded such value-addition to PC and array companies. Perhaps no drive company presently has the software, user interface, and marketing expertise to pull off a truly friendly single-board server in a brick, but all the majors are now attempting to move toward higher value added through downstream acquisitions. For example, Quantum's Meridian subsidiary has a $500 stand-alone server with multi-platform networking. We can foresee that a firm which does go after markets like this aggressively, and innovates genuine ways to reduce the cost of using digital data to do useful work, will capture significant profits. Because there are so many niches of end uses and corporate computer architectures, this market is potentially very fragmented, making it a good candidate for market segmentation, and consequently making these profits more sustainable.

Intel's Pentium: Market Meaningless Improvements

The personal-computer market continues to mature. Intel's "Intel Inside" campaign shifted the microprocessor from an invisible infrastructure component, to a consumer brand name, at a cost on the order of $700 million. The latest roll-out, the Pentium III, was budgeted at $300 million through a single ad agency. Yet for the vast majority of users and applications, the performance of the various Intel microprocessors are now effectively indistinguishable. This is a dirty little secret of Intel and AMD, but it's an opportunity for other PC subsystem makers, including HDDs.

Studies of computer usage over many platforms and applications show that speed improvements of less than 20 percent are invisible to most users, while improvements below 10 percent are invisible to the naked eye, even by experts. Notice that the measure is in percent, not in absolutes like megahertz or seconds, since human-sensory perception is logarithmic. Thus the difference between a 600-MHz Pentium III and a 550 MHz Pentium III is invisible to most of us. Furthermore, most applications are constrained primarily by the speed of other subsystems, especially networks, HDDs, main memory, and for games, the video card. Yet the price difference between 600-MHz and 550-MHz Pentiums was $200 in October 1999. For half this amount, users would get a much more significant speed improvement by buying a faster HDD or another 64MB of main memory.

Despite the irrelevance of small speed deltas, Intel always prices the latest processor at a substantial premium. Figure 1 shows selected processor speeds versus cost, with speeds

shown as relative rather than absolute levels. The 25 percent improvement in speed costs $238. In fact, there are more than 10 products covering about a 2:1 range of performance. Why bother with all this undetectable differentiation, and why such a steep price gradient at the top end? Primarily because this encourages and takes advantage of users' misbelief that small differences in megahertz matter. People, including professional reviewers, seem to measure speedups as arithmetic changes rather than logarithmic ones. They say "This machine is 50 MHz faster," rather than "This machine's CPU speed is nine percent faster." By differentially pricing small speedups at the top end, Intel is basically price discriminating to extract the maximum producer's surplus. The targets are the uninformed or testosterone-influenced buyers, e.g. those who buy the latest machine for prestige. PC makers are willing to play along with Intel and promote these meaningless differences, since it encourages end users to replace their machines faster and allows the PC makers to differentiate their products more charging an additional markup for the latest machines. It also helps that the manufacturing process for microprocessors inherently creates a spread of processor speeds. So there is little additional cost in producing this detailed product variety for Intel, although there certainly is for the PC makers.

Both meaningful and meaningless product differentiation increases in most industries as they mature toward mass consumer markets. Meaningless differentiation is especially common when products are bought for prestige or self-image. Look at the success of four-wheel drive vehicles in Southern California, where it never snows and 95 percent of the vehicles are never taken off-road. I believe that we should look to consumer industries for predictions of the evolution of many "high-tech" markets. Intel is well aware that its microprocessors are hybrids between consumer and industrial products.

Most computer component industries continue to think of themselves as making industrial products, with almost no brand recognition or product differentiation in the eyes of purchasers. In part, they may feel they can't afford consumer advertising, although massive individual advertising is not the only way to market. But also, I think some executives and product developers have trouble shifting their mindset away from the known and trusted paradigms: "If you make more capacity, they will want it."

Hard disk drive makers could do a much better job of product differentiation in the eyes of consumers. The discussion of desktop RAID looked at one approach: differentiate via downstream software and storage system design, to make end users' lives significantly easier. Even at the raw drive level, there are many differences among drives of the same capacity: interface, rotational speed, acoustic noise, power consumption, etc. The better PC reviewers realize that the drive matters to total system performance in most uses. Here is the assessment of one guide to drive purchases:

The two specifications you should care most about are two that you'll rarely see advertised: average sustained transfer rate and average access time. All the performance specifications you see most often—burst transfer rate, internal transfer rate, rotation speed, average seek time, latency, and density of data on the disk—matter only because they affect these two key specifications.

This is a failing of the HDD industry in that it has not promulgated, certified, nor marketed some important and meaningful performance differences among drives. As Intel shows, even meaningless differences can be sold to some customers. With so much confusion and downright ignorance about drive characteristics, it's no wonder that there is little consumer willingness to pay for differentiated products.

To compete effectively, through this kind of product differentiation, requires both collective and individual action by HDD makers. Increasing awareness of product characteristics, such as access time and acoustic noise, is most credibly and cheaply done collectively. It requires neutral standards for meaningful drive

performance, and educating reviewers to ask about, measure, compare, and explain these characteristics in their reviews. IDEMA, or a similar organization, is a logical, centralized location for these activities. It would be greatly assisted if one or two major PC makers would push these attributes as a way of differentiating their own products. Individual drive makers would then be able to market the superior performance of their products, with respect to specific characteristics.

Conclusion

I see both opportunities and needs for increased product differentiation, increased storage subsystem convenience to end-users, and increased end user awareness of same. Both substantive and cosmetic differentiation can be used. There are three major benefits of movement in this direction. First, by increasing the number of dimensions along which firms can compete (beyond price and capacity), it allows for more market niches and less pressure for pure price competition. Second, it increases both the true and the perceived value of the storage subsystem to people who make the PC-buying decisions. This gives storage makers a larger share of the total PC budget, at the

expense of Intel and other subsystems. Third, it increases true value to end users by providing better products, more closely matched to the particular needs NOT end-users, for different kinds of PC performance.

Of course, progressing along these lines is not easy. Some problems are internal, such as the difficulties of integrating software and systems expertise. The shift from a pure technology focus to a marketing focus on how to make end-users' lives better is also very difficult for high-tech firms. And no matter how well done, industry structure will continue to be a factor. With many competitors and short lead times for new products and capacity, successful innovations in different directions can be imitated. Nonetheless, it won't pay for each competitor to get good in each dimension of performance. With less head-to-head competition, there will be less pressure to compete purely by cutting prices.

Roger Bohn is an associate professor at UCSD. He specializes in the management of technology. He is also the director of the Information Storage Industry Center (ISIC), which does research on the business side of the storage industry. ISIC publishes a free weekly

References

1. *This article presents my personal opinions. My thanks to many for contributing to these ideas, and especially to the IDEMA Board of Directors and to Martin Kenney. Russell Stern asked the questions that got me started. All mistakes are my own; let me know what they are.*

2. *As speeds increase faster than the complexity of applications, even logarithmic changes become less important in terms of time saved. When it took 10 seconds to recalculate a 200-cell spreadsheet, cutting that in half saved five seconds. But today going from 1 to 0.5 seconds saves only .5 seconds, which is much less important. To reiterate, only applications that push the limits of CPU power benefit from CPU speedups. Aside from games, video, and engineering workstations, there are few such applications in common use.*

3. *To be fair to PC makers, they are also speeding up other system components roughly in parallel with the CPU, and consumers are more used to hearing speeds measured in processor megahertz than any other way. But end users would get more bang for a given buck if they bought the faster machine with a less than top-end microprocessor.*

4. *Note that the semiconductor industry has traditionally viewed this speed spread as an expensive liability. It took marketing genius to turn it into an asset.*
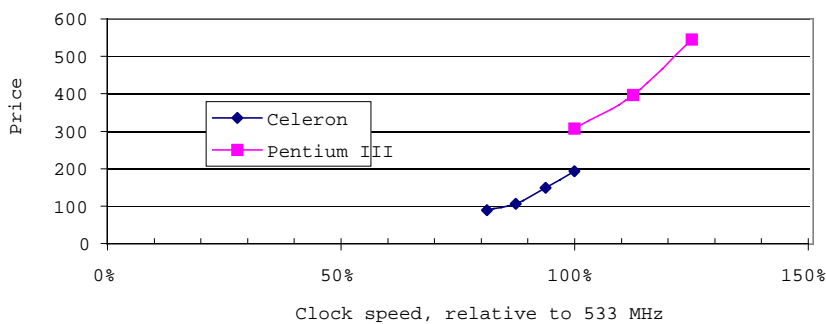

Figure 1. Pentium Processor Speed Vs. Cost