Information-Based Strategy:
New Patterns of Power and Profit

Final Examination

3 May 2018

NOTE: The exam questions total 105 points.

This examination is open book and open notes. You may refer to texts, class notes, your homework assignments, course readings, and anything relevant you find on the internet, but you are not permitted to discuss the examination with other students or former students, orally, or by text messaging, email, or other means. Please read all the questions before starting, determine where your time can be most effectively used, and budget your time carefully.

Please note that it is **not** necessary to answer all the questions in sections I and II; you need to answer two questions out of four on section I and three questions out of six on section II.

There is a small extra credit question in Part I that is worth 10 points.

(60%) Part One — Information Economics

*Answer any **two** (**2**) of the following questions. (Answer all parts of the two questions that you have selected.)*

I. Fake News

We have seen the explosion of fake news, have seen it covered as news, and have even seen fake news about fake news itself. But rumors have been around as long as human speech, and biased, incorrect, or inaccurate news stories have been around for as long as journalism.

1. Why do fake news stories resonate as effectively as they do today? Why are they so often believed by their readers?

It is easier to target fake news so that readers receive stories that correlate with their pre-existing beliefs, and are supported by those beliefs. This gives fake new stories added credibility.

2. Why is the rebuttal of fake news as complicated as it appears to be?

Again, fake news is consistent with people's existing beliefs and are designed to be simple and credible, but are not constrained by either facts or theory. Refuting fake news may require obtaining additional facts and presenting more complex theoretical responses, especially responses that are not consistent with the readers' preexisting beliefs.

3. The framers of the American Constitution assumed that only individuals and organizations of substance (whatever that meant to them at the time) would have access to printing presses. How is the creation and distribution of fake news different today from previous centuries?

First, of course, it doesn't require a printing press. Now everyone has access to a keyboard, and everyone's posts look pretty much the same. So it is difficult to use appearance alone as a signal of quality or of legitimacy.

4. As we saw in class with the discussion of Illyrian elections, political preferences are multi-dimensional and not uniformly distributed in voter attribute space. How does that make targeting fake news more complex?

There are many aspects to our beliefs, and to be truly effective, fake news should resonate with each reader's pre-existing beliefs and set of accepted facts, and thus should appear to be true. You need to know what I know in order to avoid being caught in obvious lies. You need to know what I want to believe to get me to change my beliefs on the basis of a short post.

5. Why does well-targeted fake news travel so quickly?

There are at least two reasons. First, it's plausible. Second, it's exciting, and seems worth sharing. Moreover, the guys who did the initial post are also very good at getting an army of bots to "like" and to forward their fake news, giving it the appearance of broad support.

6. What is the role of Facebook in the facilitation of targeting fake news?

Facebook makes it easier for creators of fake news to know how to target their fake news, makes it easier for them to post their fake news, and makes it easier to get likes and make fake news appear accepted and appear legitimate.

7. Why is targeting fake news so profitable for Facebook and so consistent with its desire to grow traffic and traffic-related revenue?

Again, there are at least two reasons. First, Facebook can charge for lists of people with sets of beliefs. Second, fake news is exciting, and builds online participation, which makes Facebook's sales of lists more valuable when the people on them are more actively engaged with Facebook.

8. Mark Zuckerberg has acknowledged that some regulation of Facebook may be necessary. If you were Mark Zuckerberg, what regulation would you propose?

As little as possible! I might try to limit bot access, limit bot's ability to "like" or promote content. I might even be willing to identify which posts were from outside the US, if I believed I could do that accurately.

<<Note to graders: I don't have a monopoly on cleverness here. I would accept almost anything that students offered as long as it was minimal, and created the appearance that Facebook was limiting the most obvious abuses.>>

9. If you were the head of the FCC (Federal Communications Commission) or FTC (Federal Trade Commission), what would your objectives for regulation be?

<<Note to graders: This is complicated, and a wide range of answers will be fully acceptable.>>

As the head of the FCC I cannot limit what people say, except for some few accepted restrictions on limiting hate speech and on limiting inciting violence or inciting criminal behavior. But I would want to limit the power and the impact of manipulation of public opinion that was based on outright lies. It is easier and easier to do real time fact checking. So I guess I would try to allow all speech that was not forbidden, but make it

easier to assess the quality of anything that was posted. We might assess the truth of claims that were made when assessment was possible, and report the historical accuracy of the author of this post in the past, and of accuracy of the sources references by this post in the past.

10. We did not need regulation of communications companies before AT&T. We did not need regulation of radio before AT&T launched the first advertising-based radio network. Clearly we did not need regulation of social networks before. Why might we need them now?

Again, my principal concern is that highly manipulative fake news can be created and promoted from outside the US. But I am equally concerned that the boundaries between journalism and entertainment, and between fact and opinion, and between opinion and manipulation are all blurred because there are no longer obvious signals of quality and integrity; there is no longer any cost to develop fake content that can masquerade as news.

II. Signaling, Screening, and Information Asymmetry

Information, information asymmetry, and information advantage have been central concepts throughout this course.

1. Akerlof has explained how bad sellers drive good sellers out of the market for used cars, creating the market for *lemons*, which has the value of all cars in the market collapsing so that only the worst used cars remain in the market. You can assume that the worst cars are worth far less than the best cars in the market. How would this process be altered if buyers place a value on the used cars that was 50% higher than the value that each seller placed on his or her own car? Would the market still collapse? If so, what is the final value that would emerge for used cars in the market?

Yes, the market would still collapse. However, it would not collapse to zero, or even to the lowest value of cars in the market. It would collapse to 1.5 times the value of the worst cars.

2. Is there a similar *lemons* problem in health care? If insurers were required to offer insurance to all applicants at the same price, but were free to set their prices, what would you expect to see emerge in the market for health insurance?

Prices would rise to the cost to serve the worst customers in the market. [This might be altered slightly, given that people exhibit risk aversion].

3. What if insurance companies were required to offer insurance, priced at their average cost to serve all customers, but were not required to offer insurance to all applicants. What would the market look like?

Then the worst customers would be denied insurance [while the best customers might refuse to buy insurance because they felt it was over priced].

4. What if insurance companies offered significant discounts to applicants who volunteered to wear tracking devices *and* whose tracking devices showed a healthy life style with plenty of exercise and low exposure to alcohol and nicotine. Is the company engaging in data mining? Are customers sending a signal? Is the offer of the tracking device a screening mechanism?

Yes, yes, and yes. The company is data mining to learn about its customers. Customers who refuse the device are assumed to be worse, so the company is offering the device as a screen. Customers who accept the device and change their lifestyles accordingly are working to send a signal.

5. Garces Trading Company has an attractive all-you-can-eat steak and fries menu on Tuesdays. Some of us can only eat one steak at a time. Some Warthogs Rugby players can eat several. As Pennsylvania Governor Ed Rendell has pointed out, no one requires him to invest in an individual dinner mandate, requiring him to buy an all-you-can-eat steak dinner from Garces if he is not hungry. He wants to know why it is necessary to have an individual mandate requiring us to buy health insurance if we are not sick. He likes the idea of insurance based on the average cost to serve all people, and to have it available for purchase at that low price any time you decide that you actually need it. What would you tell him about the similarities or the differences between health care and steak care?

The cost to serve differences in an all-you-can-eat steak restaurant might be a ratio of 5 to 1 between average and the extremes at Garces. And there are very few 10-steak diners. The difference between average cost to serve at Garces, lower decile, and highest decile is actually quite small. And very few of us can delay paying for dinner until we become 10-steak diners. In contrast, the difference in medical costs between average and the best decile or the worst decile is much greater. And many of us could indeed afford to cover our own expenses out of pocket until we suddenly because worst decile, and at that point we could all leap in and buy health insurance at an unreasonably low price. The system would collapse.

6. Describe a signal or screen you designed and used in your own life. Did it work the way you expected? If so, explain why, and if not, explain why.

When I rented an apartment in Ithaca New York I offered to leave a full year's security deposit, as an indication that I truly was not concerned that my young daughter would damage the owner's property. The signal was accepted and we were allowed to rent the house.

<<Note to graders: Accept almost anything as long as (1) it was intended as a signal or a screen and (2) it is properly described; that is, if the intent of the designer of the process to communicate about *himself* then it is a signal and if the intent was to learn about the *other party* then it is a screen. Finally, (3) did it work, and does the student know why it worked?

7. Why are screens preferred by corporations, which rarely rely on signaling except in some cases of hiring from top universities?

Because it is so difficult to calibrate a signal, to know how strong a signal is, or to guess how hard it would be for some individuals to fake a signal.

III. Valuing Investments Under Uncertainty

You and the (hardworking) members of your study team have yet another plan to get rich quickly, this time by developing a hyper-secure form of electronic communication. You've developed *SnapBabble.* You dictate your message and it is stored on your phone. Your phone automatically sends a text message to your recipient; when the recipient responds your recording is played, once, and then deleted. No record anywhere. OK, your recipient might find a way to record the message, but, hey, nothing is perfect. Once again, you and your team will be wealthy.

As always, your marketing research suggests that there are three possible ***outcomes***, each with known payoffs and known probability of occurring. As with so many software ventures, your payoff will vary with the time you enter the market. If you are an early mover *and* there is a market for your product, your payoff will be greater than if you are a late entrant. Unfortunately, there is a third possibility, that of being semi-late; if you follow the Phased strategy you studied previously you will be neither early nor late but somewhere in between.

Outcome	**Probability**	**Early Payoff**	**Semi-Late Payoff**	**Late Payoff**
No Demand	1/3	Z0	Z0	Z0
Medium Demand	1/3	Z4500	*Does Not Apply*	Z2100
High Demand	1/3	Z9000	Z8100	Z4200

As always, there are multiple ***development strategies***, which will have different costs, and will have different timings for completion, which in turn will have different implications for your payoffs.

Assume you have three different development strategies:

- **All At Once (AAO)** — You spend Z2400 immediately to build a Large system, and you can handle whatever demand you get.
- **Phased** — You spend Z1800 up front, to build a Small system, which is enough to handle Medium Demand. It is not enough to handle High Demand. If you spend an additional Z750 you can handle High Demand as well, and you will receive the Semi-Late entrant payoffs for High Demand.
- **Deferred** — You spend nothing until you see what the market looks like! You can build a Small system for Z1800, which will handle Medium Demand. You can also build a large system for Z2400, which will handle High Demand. In this case will receive late payoffs for the demand levels you build to accommodate.

1. First, you need to know what your strategy should be. Suppose you choose **Phased** implementation. When if ever should you make the second investment and build a Large System? Suppose you choose to follow the **Deferred** implementation. When should you build nothing? When if ever should you build a Small System? When if ever should you build a Large System?

With **Phased** I always start by building small only upgrade to large when I encounter high demand. With **Deferred** I build small if I encounter Medium demand and I build large if I encounter high demand; otherwise I do nothing.

2. OK, now you know what you are going to do under each of the three development strategies. Right now, before you know anything else about the market, what is the expected value of each of the three strategies? Which is the best?

Dev Strategy

	AAO	Probability	Payoff	Exp Payoff		Exp Cost
Outcome	None	0.333	0	0		800
	Mid	0.333	4500	1500		800
	High	0.333	9000	3000		800
	Total Payoff			4500		
	Total Cost			2400		
	Net			2100		

Outcome	**Deferred**	Probability	Payoff	Exp Payoff	Incurred Cost	Exp Cost
	None	0.333	0	0	0	0
	Mid	0.333	2100	700	1800	600
	High	0.333	4200	1400	2400	800
	Total Payoff			2100		
	Total Cost			1400		
	Net			700		

Outcome	**Phased**	Probability	Payoff	Exp Payoff	Incurred Cost	Exp Cost
	None	0.333	0	0	1800	600
	Mid	0.333	4500	1500	1800	600
	High	0.333	8100	2700	2550	850
	Total Payoff			4200		
	Total Cost			2050		
	Net			2150		

Clearly, **Phased** is best.

3. Suppose once again that you had access to the *Insider Trading Genie*, who could tell you in advance what demand would be? What would your ideal strategy be? What would your costs be under the ideal strategy with perfect information? How can you know, without doing any calculations? What would your payoff be? How can you know, without doing any calculations? What is your net payoff with perfect information? What, therefore, is the value of perfect information?

With perfect information I would incur the implementation costs of **Deferred** and the Payoffs of **AAO**. That's a payoff of Z4500 and a cost of Z1400. That's a net payoff of Z3100, or Z950 better than my best alternative. The value of perfect information is therefore 950.

4. The **Phased** strategy can sometimes be dangerous, because your starting can be observed by others. Suppose Bringle has decided that it likes *SnapBabble* and that it wants to integrate it into all Android Phones. Suppose that your market intelligence unit concludes before you start your calculations that if you follow the

Phased implementation *and* if there is High demand Bringle will notice you. Suppose that in that case there is then a 20% chance that they will ignore you, leaving your payoffs unchanged; there is a 20% change they will destroy you by integrating their own version, *BBabble*, directly into the Mannequin Operating System, dropping your payoffs to Z0; and a 60% chance that they will buy you out for Z5000. If they buy you out for Z5000 that is the only payment you will receive. You will not receive any of the payoffs based on serving demand. Fortunately, they announce their decision before you have made your second investment. What small part of your analysis changes? (For this and the following question you no longer have perfect information.)

Nothing changes in the analyses of **AAO** or **Deferred**. But the high outcome in **Phased** would now be less valuable than it had been. I get reduced payoffs from high demand, as shown below. But my expected costs are changed, because I only make the second investment (Z750) 20% of the time that I experience high demand, the 20% of the time that Google ignores me.

Danger	Probability	Payoff	Exp Payoff	
New Payoff Calcs	0.2	0	0	<<destroy>>
For High Demand	0.6	5000	3000	<<buy out>>
	0.2	8100	1620	<<ignore>>
			4620	

5. What is your optimal strategy now?

Clearly, my best choice is no longer **Phased**. I would revert to **AAO**. The incurred cost for the High outcome of **Phased** is the cost of the of the small system, plus the additional cost of the upgrade to a large system (Z750) the 20% of the time that high occurs and Bringle ignores me.

	Phased with Danger	Probability	Payoff	Exp Payoff	Incurred Cost	Exp Cost
Outcome	None	0.333	0	0	1800	600
	Mid	0.333	4500	1500	1800	600
	High	0.333	4620	1540	1950	650

Total Payoff			3040		
Total Cost			1850		
Net			1190		

6. Suppose the payoffs for being late have changed. Suppose the payoff for Medium demand is always Z4500, regardless of timing. Suppose the payoffs for High Demand with semi-late and late implementation are both Z8100. What would you expect your best strategy to be? What's the simplest way to show this? Is there a general principle here?

You could just plug new numbers into your spreadsheet. Alternatively, you could note that the cost of **Deferred** has not changed, but that the payoffs of **Deferred** when you get medium demand are up Z2400 and the payoffs from high demand are up Z3900. Each occurs one third of the time, for an expected increase of Z2100. If I add Z2100 to the value of **Deferred** the value is now Z2800, which makes it my best choice. The general principle is that when there is very little cost to waiting until I can act with perfect information, waiting becomes the best choice.

	Alt Deferred	Probability	Payoff	Exp Payoff	Incurred Cost	Exp Cost
Outcome	None	0.333	0	0	0	0
	Mid	0.333	4500	1500	1800	600
	High	0.333	8100	2700	2400	800
	Total Payoff			4200		
	Total Cost			1400		
	Net			2800		

IV. New Business Models and Regulatory Complexity

New online business models often have interesting regulatory complexities that were not fully anticipated with previous business models.

1. Two-part tariffs are a great mechanism for a monopolist to deploy to protect and extend its monopoly. When Microsoft changed its pricing to manufacturers like Dell it changed from a high cost per copy of Windows, and instead charged a

fixed cost per year plus a very low cost per copy. The two parts of the pricing were set so that Dell got a slight savings, and so that the total reduction in Microsoft revenues was small. Why did this make it almost impossible for a new entrant to compete with Windows? Why is this two part tariff so much easier to implement with software products than with physical goods?

Using numbers helps make this clear. When Dell paid $50 / copy then competitors merely needed to come in below $50 to have a chance to sell. When Dell paid a large fixed cost like $45 million annually and $2 for each of a million copies Microsoft gave up a little and dell saved a little, but now competitors had no way to compete. [You can offer $35 to compete with $50. How do you compete with company that charges only $2 for each additional copy[, especially when Microsoft is still receiving close to $50 per copy when both parts of the price are considered]!

2. Third party payer gateway systems like search often have monopoly power even when they are not monopolies. Why?

This comes back to single homing and mandatory participation. If a searcher uses only one search engine, then a seller who is not listed in both gateways systems loses the customers who use whichever system the seller no longer uses. Since the seller needs to be in both gateway systems, the two gateways really do not compete on the prices they charge the seller.

3. When third party payer gateways systems like search have monopoly power they often engage in reverse price wars. Why are reverse price wars seldom seen outside these gateway systems?

It makes sense for the gateway operator to charge the sellers more, and thus to pay more to searchers, and thus to become more important to sellers. But this only works when there are parties that use the gateway to find sellers, and sellers who use the gateways to be found by buyers. If search charged searchers more, they would leave and use competitors' search platforms. If search charged sellers more, they would leave unless the gateway had become essential to their being found by searchers. So these really are the only places where you see reverse price wars.

4. Platform envelopment is yet another great way to extend a monopoly. Google's Android platform seamlessly integrates a wide range of offerings, and uses the MDSA to tightly delimit what other companies can have pre-installed on Android devices. Why do platform envelopment strategies often appear to be so beneficial for consumers? Why do some regulators object to platform envelopment strategies?

Because the parts of a platform interact so seamlessly, they create super-additive value. [Additionally, there is often one part where it is difficult for the consumers to determine what the price could be for the item alone, which means that everything else can be added, apparently "for free."] Regulators object because platform envelopment is often used to kill competitors, either by denying them access to the platform, or by offering the platform's own products apparently free.

5. The *"myth of anonymity"* says that since Google does not actually read your mail and does not actually give your email address to advertisers, you are safe no matter what they know about you. Google merely sends you ads on behalf of advertisers who requested someone with exactly your characteristics. How can this harm you?

First you need to understand that sellers can use information about you to alter what they offer you or what they withhold from you, and to alter the prices they charge you. Then you need to understand that sellers know exactly what combination of conditions they used to define the intended recipients of a targeted ad. So even though a seller is not told in advance who you are, the seller knows what they should offer you and what they should charge you as soon as you respond, because they know which ad you are responding to. This can, and already does, result in your paying higher prices when the seller can anticipate that you are willing to pay higher prices.

6. The "right to be forgotten online" is a proposal, enacted into law in some markets, that allows users to ensure that some information about them cannot be retrieved through search. Information remains on the net, but in theory it cannot be found by someone simply by searching for online references to you. Why does this not mitigate the risks of (5) above?

The right to be forgotten online has nothing to do with Google's ability to characterize you for sellers wanting to know about you so that they can send you targeted ads. [Something can be removed from search results, but that does not affect what is available online or what Google knows about you and is allowed to use about you.] [Additionally, most things that are removed under the right to be forgotten are old and the user can safely argue they are no longer relevant. Old and irrelevant material is not the basis of Google's targeted ad campaigns.]

7. Why, then, would the right to be forgotten online have no impact on Google's business model or Facebook's business model, as they exist today?

Google's revenues do not come from organic searches about individuals! Google's revenues come from profiling, advertising, and keyword auctions. None of this is affected by the right to be forgotten.

8. As we discussed in class, the Fairness Doctrine is seldom applicable today, except when *all* traditional media are covering an event like *The State of the Union Address*. Surely, anyone can post on Facebook. What is the basis of the argument that Facebook has an obligation to limit the propagation of fake news, that is, that Facebook has to be in some sense an arbiter that limits free speech?

The idea is that fake news is carefully targeted to resonate, and that fake news is difficult to distinguish from real news because it appears to be as legitimate as any other news feed and it is targeted to resonate with its individual readers. If fake news were not so very effective, then Fakebook's argument would probably need to be taken more seriously.

V. Framing (10 Extra Credit Points)

Extra Credit: A bag contains one counter, known to be either white or black. A white counter is put in, the bag shaken, and a counter drawn out, which proves to be white. What is the chance of drawing a white counter now? Explain how you got your answer.

There is a 2/3 chance of drawing another white counter.

Combinatorial analysis 1:

After I put in a white counter, there are two possibilities, white/white and white/back. Pulling out a white counter is twice as likely with W/W than with W/B, so the fact that I pulled a W means that I am twice as likely to have started with W/W.

Combinatorial analysis 2:

Imagine I put in a green counter and then I pull out a not-black counter. If there had been a white in the bag initially and I added a green, then there are two ways to pull out a not-black. If there had been a black, then there would be only one way to pull out a not-black. So, again, after the fact does tell me something about what I started with.

(45%) II. Short Answers

*Answer any **three (3)** of the following:*

A. Scenarios

I know how to use sensitivity analysis, to see how differences in parameter values affect the overall value of a plan. (i) What does Monty Carlo analysis tell us that we can't learn from deterministic sensitivity analysis? (ii) How is Scenario Analysis different from Monty Carlo analysis or sensitivity analysis? (iii) What is meaning of the long phrase, *"the things you can't know, such that if you could know, you then would know, exactly*

what you needed to know"? (iv) Given that you can't know those things, and therefore don't have the data, how is scenario analysis useful at all?

(i) Monte Carlo analysis allows us to place probabilistic bounds on our worst case expectations. It is not enough to say that on average we need the following line of credit. We also need to know what we might expect the worst 5% or the worst 10% of the time.

(ii) Monte Carlo analysis and sensitivity analysis both allow us to explore *different values for parameters*. Scenario analysis doesn't change the value of a *parameter*; it *changes the entire story*.

(iii) These are the most important things; if you could know them you would know your answer. But you can't know them. [They are you drivers, or your critical uncertainties.]

(iv) It helps you impose a structure on your uncertainties, which allows you to analyze a sequence of related problems[, rather than simply be overwhelmed by uncertainty].

B. Fair vs. Efficient Marketing and Regulation

(i) Is it fair to charge all applicants the same price for health care? (ii) Is it fair to let Uber drivers to compete with licensed taxis? (iii) Is it efficient to charge higher insurance prices for sick people, or for people with genetic predispositions to severe illnesses? (iv) Is it efficient to allow Airbnb rentals to transform neighborhoods and displace people from the rental apartments at the end of their leases? (v) Why are questions about fairness and efficiency arising more often today than in previous decades?

(i) I don't know what *fair* means here. It's *good* for sick people and *bad* for healthy people.

(ii) I don't know what *fair* means here. It's *good* for passengers and for Uber drivers, *bad* for people who purchased taxi medallions and paid for commercial drivers' licenses and insurance.

(iii) Yes.

(iv) And yes.

(v) You didn't need to worry about the implications of perfect pricing on insurance before increased information availability made perfect pricing possible. You didn't need to worry about competition between professionals and participants in the

sharing economy, or about the implications of the sharing economy on others, before online spot markets made the sharing economy possible.

C. Resonance Marketing

Resonance marketing is a new strategy, a form of newly vulnerable markets for niche products. (i) What was the role of the uncertainty discount in preventing these niches from being exploited by sellers? (ii) What is the role of *informedness* in allowing sellers to enter these markets? (iii) Why is resonance marketing most effective when selling to high V high t customers? (iv) And why was the uncertainty discount most effective in blocking sales to high V high t customers?

(i) The uncertainty discount said that if consumers don't know what something really is, they value it as the average of all of the things they thought it might be. Niche products appeal to consumers who believe the product might be perfect for them, but until the uncertainty discount could be reduced, these consumers were unwilling to search out and purchase these niche products.

(ii) Until *informedness* reduced or eliminated the uncertainty discount, the cost of advertising was too high for niche products and producers could not reduce the uncertainty discount sufficiently to enter the market.

(iii) Niche products and resonance products generally are more expensive than mass market offerings. They appeal to customers willing to pay more to get exactly what they want, which is of course the high V high t market.

(iv) Because high t customers have the highest uncertainty discount[, since for the same range in product attribute space the higher t creates a higher compromise discount.] [This matters because the uncertainty discount is just the average of the possible compromise discounts.]

D. Newly Vulnerable Markets and the Difficulty of Responding

When Capital One started to capture AT&T Credit Card's best customers AT&T saw only two alternatives. They could reduce the APR for *all* of their customers *now* and *keep them all* but at lower profitability. Or they could reduce the APR for those customers they were about to lose and *hope to retain some*, again at lower profitability. (i) How would you choose between the two alternatives? (ii) What data would you need? (iii) How would you respond to an executive who said that both alternatives left you worse off than you are now?

(i) I would need to determine which lost me more money quickly and which cost me more customers over the long term.

(ii) To do that I would need to know how much money I was earning from the customers I was losing and the speed with which I was losing them. I would need to know what interest rate I would have to offer them immediately in order to keep them and how that would affect my earnings. I would also need to know what interest rate I would have to offer to retain those customers that I was about to lose, and what fraction of them I could retain if I matched their offers from Capital One. [This would allow me to evaluate the immediate collapse of earnings due to lower interest rates on my existing customers; this at least would remain stable since I would not be losing customers. I would need to compare this immediate loss against the gradual but potentially more complete loss of customers and of earnings if I lost customers who found Capital One's offer attractive, and if my attempt to retain customers were less than totally effective. I would not lose profitability as quickly, but ultimately my profitability would suffer more.

(iii) I would say that was true but irrelevant. The question is not whether I would have been happier without Capital One as a competitor, but rather what is the best I can do now that they have arrived as a competitor. [I could also simply explain the trap of the wrong base case.] My base case is probably now the scenario in which I do nothing, which is probably worse than either of the alternatives considered in (ii) above.

E. Newly Vulnerable Online Markets and Channel Power

At the time that airlines started to sell air travel online, consumer companies like Lever and J&J considered selling online as well. (i) How were companies like Wal-Mart and Home Depot able to persuade their suppliers not to consider online sales direct to individual consumers? (ii) Now consider services where adoption of new alternatives is slow, like insurance, or where customers are loyal to their account executives, like investment advising. How would you expect companies in these industries to approach online sales and service? (iii) Now consider *inspection goods*, items like very expensive Patek Philippe watches, which consumers are unlikely to purchase unless they are able to examine them and see how the feel and how they look on their own wrists. How important are retailers to these manufacturers? How would this affect the relationship between manufacturers and retailers?

(i) Suppliers were told that they would be dropped from the stores, and that their sales would collapse as a result. [They were told that the collapse would be immediate, and that consumers would move to online sales only slowly if at all.]

(ii) Insurance companies and investment advising companies were very slow to adopt online sales, because they were afraid of alienating their best agents / account executives.

(iii) [Cartier owns all the little Cartier boutiques you see embedded in department stores, so that indeed online sales if they existed would be competing only with Cartier.] Patek does not sell watches online! Patek will not risk being seen as competing with its essential retailers and their show rooms. Patek will offer to put you in touch with one of their retailers.