[*Foundations of Interaction Design:* Ch. 18 Picci: 1st sub-edit, 23 Feb 2005, by Philip Tabor]

[Body text and associated headings: 11 pp., 4,103 words, 25,535 characters (with spaces)]]

Chapter 18

Bad Design by Design? Economics Meets Other Types of Interactions

Lucio Picci

1. ECONOMICS AND DESIGNING: THEIR COMMON GROUND

A few centuries ago the discipline of economics was not yet a well identified and separate branch of knowledge. Its foundation and its subsequent development required hardships and toil by a group of pioneers coming from disparate fields of knowledge such as moral philosophy, physics, and engineering. Their efforts have produced a rigorous theoretical construction that many consider as a true sanctuary, providing its priests – the economists – with a set of powerful tools to analyze reality. As an economist myself, I share a sense of respect for the intellectual construction, and its foundations, that our predecessors endowed to us.

In the eyes of many detractors of the science of economics, indeed, the project proved to be too successful, because it emboldened economists to the point of arrogance, and they started using their analytical weapons to conquer and dominate bordering fields of knowledge. Other lands of knowledge – such as the political sciences, sociology and geography – have become fair game for economists, who are now seen by their indigenous inhabitants as an invading army with no United Nations' approval: economics is sometimes seen as an imperialistic social science, and it is treated with an equal share of scorn, respect, and envy.

So before any misconceptions arise between the indigenous inhabitants of the honorable design discipline and myself, let me assure you that I come in peace, and that my only purpose is to contribute to a fruitful dialogue between our two seemingly very different fields of knowledge (and maybe to distract you as our troops get ready).

As I will argue, this difference is mostly one of looks, and may just follow previous habits, because designing and economics increasingly share a common ground. In particular, both have to do with describing interaction systems, an endeavor where economists have learned lessons that are often subtle and surprising. For example, sometimes the optimal design of a product is not the one that provides the interested parties with the best possible interaction experience, but is an intentionally "bad" design. Such an outcome may seem illogical to all those designers whose mission is to do their very best in the service of the final user of a product. Contrary to that belief, we sometimes observe examples of "bad design by design."

Another issue that I raise here is the organization of the design, prototyping, and production process, seen in the light of open-source software production, an emergent and interesting designing and production method. I will draw some parallels between open-source software production and the design and prototyping process familiar within the interaction design community, in order to discuss their similarities and their differences.

As a prerequisite for both discussions we need to understand in what sense economists increasingly consider themselves as designers and engineers.

2. "MECHANISM DESIGN": THE ECONOMIST AS AN ENGINEER AND AS A DESIGNER

The objects that economists design are called "mechanisms," a set of rules to solve a problem of allocation (of goods, other resources, or people) in a situation where there are conflicting goals. A mechanism is an interaction system in the sense that it gives form to, and orientates, a set of interactions between different economic agents. An example clarifies the point.

Assume that Interaction Design Institute Ivrea organizes an internship program for its students. When the numbers are not too big, this type of task can be effectively carried out informally, trying to juggle all the requests coming from the different actors – the students, and the organizations who participate in the program. However, many conflicting goals are present, and the allocation problem of an internship program is conceptually challenging. Some students are better than others, and so are some of the organizations. The students would like to do their internship within a good organization, but firms also would like to have the best students,

particularly when they look at them as potential future employees. Moreover, students differ among themselves in other dimensions: They live in different places and may prefer an internship not too far from home. Some, moreover, have spouses and possibly children, so their degrees of freedom are limited by these facts of life.

Guaranteeing the best match between students and organizations is not easy, and "matching theory" – a branch of economics – is devoted to understanding these types of problem. When this theory is used not just to understand but to solve a practical matching problem, by indicating a set of rules to be followed when managing it, we have an instance of "mechanism design." More details on the "internship problem," incidentally, are in Varian (2002), with reference to Roth (2002).

3. PRICE DISCRIMINATION AS MECHANISM DESIGN

The next example of mechanism design has to do with what economists call "price discrimination," and it leads us to the possibility of rationally desiring a "bad" design for a product.

Producers of goods (or services) would rather not sell at a uniform price, but prefer instead to impose a higher price on those willing to pay more for the product – because they are rich, or because they very much like it. As long as they recover (marginal) costs, producers are also more than happy to sell at a lower price to people who would not otherwise buy the product. "Price discrimination," as this activity is called, is very common. But it is not straightforward to achieve, because potential customers will not easily let the producer know their willingness to pay: people who are asked how much they would like to pay for a given product, knowing that what they declare will reflect itself in the price tag, probably understate the truth.

However, it turns out that there *are* ways to know, at least approximately, how much people are willing to pay for a product or, in economists' jargon, to obtain a "truthful revelation of preferences." This knowledge is often obtained by creating options where, by choosing according to their self-interest, people implicitly (and usually without being aware) declare their willingness to pay. It is another instance of "mechanism design": When the mechanism is designed properly,

customers end up self-selecting themselves into the appropriate category, where they are made to pay according to their true willingness to pay.

A few examples clarify this point. In the market for books we can roughly assume that there are two types of reader. Some are avid readers of a particular author's books and will pay much for them; sometimes, indeed, the publication day of a new book by a well known author amounts to a memorable day in the life of her fans (ask my nephew about Harry Potter). Other readers, by contrast, are less interested and will only buy the book at a lower price.

The publishing house wants to keep these two types of customers separate. This is achieved by distinguishing between the hardcover edition and the paperback. The fact that the former has a better binding, and is often printed with a bigger and more readable font, is in fact irrelevant: the difference in cost between the two editions is negligible and in no way justifies the difference in price. What is relevant is that the cheaper paperback edition is typically published after an intentional delay. The self-selection of customers is obtained through the difference in their impatience: The eager reader wants the book as soon as possible, while the less interested reader can wait.

Mail-in rebates are another instance of price discrimination. A rebate is offered to the customers of a product who take the time to fill in a form and send it to the producer. Those willing to pay more, who tend to be richer, value their time more than others, and will not bother to mail the rebate: They pay the full price. On the other hand, the people who do apply for the rebate, by taking the time to do so, implicitly declare their low willingness to pay for the goods, and end up paying the lower price.

One last well known example of price discrimination by self-selection is provided by the air travel industry. The price of the same seat on an airplane can vary significantly, depending on the type of ticket, on how far in advance of the flight it is bought and, more generally, on the restrictions it carries. A typical restriction involves stopping over at the destination on a Saturday night. There is no technical reason for this – airplanes are not required to rest away from home on those nights. The rationale has to do with price discrimination: tourists – the people with a low willingness to pay – do not mind, and often actually prefer, staying out on Saturdays; business

people, on the other hand, who are willing to pay more for their tickets, want to be back from their trips for the weekend.

What about “bad design”? All these examples of price discrimination, and many others, involve the introduction of some type of bad designing. Consider mail-in rebates. If the purpose of the seller is to give a rebate of, say, $2 on a $10 purchase, there is an obvious way to design the interaction with the customer better: Write “$8” on the price tag, and avoid the hassle of mailing in (and processing) the rebate. However, by choosing this hypothetical “better” design, the discount would go to everybody, including the people who would buy the product even at the higher price: there would not be price discrimination.

Paperback books are also an example of bad design: they could be improved at very little cost, and could be published promptly and without making people wait: in other words, they could be their hardcover version. Similar considerations apply to cheap, but highly restricted, airplane tickets. Without the restrictions, the interaction between customer and product would obviously be better. In all these cases, however, the inconvenience of the cheap version of the goods is for a reason: to obtain price discrimination.

Let us generalize the above examples. With mail-in rebates the bad designing is about the way a piece of goods is priced. With books and restricted airplane tickets, the bad designing also results in some inconvenience for the customer. The designing often involves not a single product but a whole suite of different versions of a product, so that they induce self-selection among different types of customers – as in first- and second-class transport in trains or airplanes, for instance. This is obtained by artificially creating some flaw in the lower price version(s) of the product, so that it is still appealing enough to the low-willingness-to-pay customers, who pay less for it, but is bad enough for the high-willingness-to-pay customers, who prefer to spend more for the better version.

In some cases, such as in the hardcover/paperback versions of books, the difference in cost to the producer between the different versions is minimal. Sometimes, indeed, the more expensive version is cheaper to produce, like the IBM Series E laser printer at the beginning of the 1990s, whose slower and lower price version was obtained by introducing a “slowing” integrated circuit into the faster and more expensive version (cited in Deneckere & McAfee, 1996). So, in order to price-discriminate, producers often actually have to design *bad* products. The job of the designer is

to design the suite of versions jointly, one in relation with the other, to obtain for the worst version(s) what we could call “optimally bad design.”

4. HOW JUSTIFIABLE, AND HOW RELEVANT, IS “BAD DESIGNING”?

There is a common perception, on whose cultural and historical origin I will not speculate, that the price of something has to be “fair.” If there is such a thing as a fair price, it must be unique, and selling the same product (or two very similar products) at different prices to different people seems unjust, unethical. So, is price discrimination “unethical”? Is this state of affairs good only for the producer, who reaps a higher profit at the expense of the customers? Would it make sense to forbid price discrimination so that, for example, we could all afford the luxury of hardcover books and unrestricted cheap flights? Not quite. Let us consider this issue from the point of view of the welfare of customers.

It turns out that price discrimination often serves the interest of not just the producer but customers too, because it allows the market to serve the people with a low willingness to pay. For example, if airlines were forced not to price-discriminate, chances are that the cheap fares at the base of most international tourism, and of some international academic projects, would not exist.

And it is not only some people who would suffer. Many industries would not exist without the possibility of price discrimination, because competition on a unique price would drive the price below the level that allows producers to recoup their fixed costs. This is particularly true for industries characterized by a cost structure in which the fixed cost for setting up the business and building the “first copy,” or prototype, of a product is high, but where additional copies of that product are inexpensive. Writing the manuscript of a book, for instance, or writing, executing and recording a piece of music, requires much effort by highly skilled workers; but printing a book or copying a CD has a very low unit cost. When the cost structure for production has such pronounced “increasing returns to scale,” as economists would say, price discrimination is indeed a crucial issue not only for the producers, in order for them to stay in the market, but also for the industry to survive and for customers to be served.

Most information goods, and many information technology products, share such a cost structure. Since these goods are increasingly important, it follows that price discrimination is a

practically very relevant issue, and the idea of having a “bad design by design” is much more than just an anomaly.

5. SERVICE DESIGN V. MECHANISM DESIGN

I have not mentioned so far that the study of organizations is one of the fields of knowledge that economists managed to colonize. For this purpose economists developed the so-called "transaction-cost theory of the firm" (and of other types of organization). According to this theory, in deciding whether to do something by interacting in a market, or within the boundaries of an organization, of paramount importance are the costs incurred in carrying out the transactions; if these costs are relatively high in an "arms-length" market relationship, there is room for setting up an organization where the same transactions are carried out by its members within the boundaries of the organization itself, and without using a market (see Williamson, 1975).

This knowledge, besides exposing my credentials as an economist in speaking about organizations, allows me to consider what interaction design could learn from a mode of organizing production that is receiving much academic attention: the production of open-source software – that is, of software distributed under a license which permits users to freely share or modify it.

It is a type of production with very little (apparent) structure, very horizontal in its organizational relationships, and allowing much room for collaboration and to experimentation. The tight relationship between production and experimentation is shown by one of the open-source community slogans, "deliver early, deliver often." There is no clear distinction between the planning and the production phase of a product.

Also, open-source developers usually do not pay much attention to the codified tenets of software engineering: the discipline that establishes how software projects should be conducted and determines, among other things, that the requirements of the software to be developed should be analyzed formally and at length. Open-source software development usually starts with the purpose of solving a problem that the developers themselves face, but does not include a proper analysis of requirements.

There are noteworthy analogies between open-source software production and some ideas familiar within the designing community. Consider the issue of governance. At least in principle, within an open-source software project the way decisions are taken does not preclude anybody's contribution. In this sense, Open-source software production is an example of what has been defined as "participatory design," a theme of some relevance within the designing community (see

chapter 10 by Pelle Ehn in this book). Another trait in common between open-source software production and interaction design is the common emphasis on prototyping (see chapter 16 by Bill Moggridge).

These are just analogies, but suggestive ones which raise issues for consideration. One issue has to do with incentive systems: What makes people willing to play the open-source game? We know that many open-source programmers work for free, and economists' analyses are based on the idea that people are self-interested, so that observing highly qualified professionals not receiving a pecuniary retribution for their services amounts to a puzzle. An explanation of this apparent conundrum is based on the observation that programmers' productivity can vary enormously depending on their skills, and that these skills are hard to communicate to potential employers. By participating in an open-source project, high quality programmers are able to reveal themselves as such, to not just their fellow programmers but the world at large. While working for free, they acquire a reputation as a good programmer, which has a significant market value (see Lerner & Tirole, 2002).

The first message that the open-source experience delivers to interaction design, then, has to do with incentives. If designers want to experiment with more participatory forms of design, they should think hard about the incentives for participation, possibly taking suggestions from the highly successful open-source community. There, persons may even work for free, as long as their good work contributes to building a good reputation which can be utilized in the future, for example in the job market. A crucial aspect of the open-source community is that individual good work does not "get stolen," to the point that one of the greatest "crimes" is to deny someone's contribution. The incentive system of open-source software production is based on acknowledging personal contributions to the project.

Another issue worth analyzing is the relationship between designing an item and its actual industrial production. This distinction is simply not present in open-source software production: the first version of a program is often made to meet the needs expressed by the programmers themselves, and is used right away, at least by the programmers, who can put up with its early idiosyncrasies and, almost always, lack of good documentation. Designing, experimenting,

producing, and using are all meshed together, and in this respect the open-source software production method itself represents an impressive interaction system at work.

Within more traditional contexts, on the other hand, the participatory part of the design process and its experimental emphasis, when present, only occur during the prototyping phase. Once the product has been designed, it changes hands and goes to the factory, which to the designer is to some extent a black box with impregnable walls. When designers have finished their job, they move on to a new project.

Open-source software production suggests new questions and a change in perspective. Could its mode of production, with such a tight integration between the designing and the production of goods, be extended to goods other than software? Could we have an open-source car, bottle-opener, or chair? Is open-source production an interesting curiosity only good enough for software, or can we have "Open-source Everywhere"?

While it is not easy to answer these questions, I find such a possibility very intriguing and worth exploring. Researchers from the interaction design world who are today thinking about how designing activity is carried out, will breath some fresh air by considering the evidence on open-source software production. A few analogies do not make a relevant case, and more research work is needed. However, we should welcome the adoption of an expanded vision of the designing *and* of the manufacturing problem, where the distinction between the two, designing and production, is not so obvious, nor so inevitable.

7. CONCLUSIONS

Economics, I have argued, is increasingly about the design of interaction systems, and its analytical tools allow new insights to the field of interaction design. I have provided a couple of example to make my point. First, I have shown a case of "mechanism design" – price discrimination – where it makes sense to produce goods that are intentionally and *prima facie* sub-optimal. Observing "bad design by design" is not an intellectual curiosity without practical relevance. On the contrary, we observe it frequently in a world where "increasing returns to scale" are ubiquitous.

"Bad design" allows for price discrimination, which should not be considered unethical, because it often serves societal needs. It follows that implementing price discrimination is in accordance with the objective of contributing, through good design, to human happiness: in order to make a good design, sometimes what is needed is a "bad" design. Foul is fair, then? Not quite: simply, things have become a little bit more complicated than they used to be, and this is healthy, because it means that we have enlarged our perspective.

I have also argued that a contribution to interaction design could come from economics as a discipline that studies organized behavior. The example of open-source software production casts light on the process of designing and, in particular, on the distinction between designing and manufacturing. True, it sounds futuristic at best to think that the open-source mode of production could be extended to other realms; that we can apply such a playful way of dealing with what we want to design, experiment, produce, and use; and that we do all this at the same time, without a clear distinction of phases. However, since the discipline of interaction design is thinking about its future, I believe that some science fiction is not out of place here.

The interaction design of tomorrow, then, should be able to consider not just the interaction between the goods it produces and the people who use them, but the whole set of actors of the system: users, producers, *and* designers. There are many complicated interactions involved, so a theory is needed to provide a framework and make the problem's complexity manageable. Economics can give a hand in this effort.

An appropriate analytical framework would allow us to see more clearly through the different issues I have presented. For such an effort, I am convinced, we would receive a double reward: Not only we would understand better, we would be able to *do* better. We would be more prepared to fulfill what I see as the highly ethical goal of interaction design: to make the whole interaction experience with the goods, material and immaterial, that enrich our lives, as enjoyable and beautiful as possible.

The author of this chapter thanks Sebastiano Bagnara, Tom Erickson and Raimondello Orsini for their helpful suggestions.

[ends]