

NEWS AND VIEWS

The economics of synthetic biology

Joachim Henkel¹ and Stephen M Maurer²

¹ TUM Business School, Technische Universität München, München, Germany and ² Goldman School of Public Policy, University of California at Berkeley, Berkeley, CA, USA

Molecular Systems Biology 5 June 2007; doi:10.1038/msb4100161

Introduction

Ten years ago, genetic engineering was limited to cutting and pasting DNA from existing organisms. Today's biologists can write down gene sequences that have never existed anywhere, place an order over the Internet, and receive the desired DNA by return mail. The new science of synthetic biology dreams of a day when blueprints for new life forms can be designed as easily as computer chips. Practitioners argue that the key is to create libraries of standard gene sequences ('parts') that reliably perform simple functions like encoding an enzyme or building a protein that detects light. This strategy is potentially powerful: the electronics industry already uses similar libraries to create ultra-complex objects like computer chips and software (Endy, 2005).

The technological benefits of introducing electronic methods into biology seem clear. The economic consequences are more ambiguous. Many electronics and software industries feature a dangerous 'winner-take-all' or 'tipping' dynamic, in which an initial frontrunner becomes steadily more entrenched over time. Microsoft's rise to power is an obvious example. Significantly, tipping dynamics do not always lead to monopoly. In fact, many outcomes are possible: The eventual winner can be open (Apache) or proprietary (Windows), technically superior (Web) or suboptimal (VHS). Historically, these outcomes have emerged more or less at random from rough-and-tumble contests in the market. Synthetic biologists can and should do better. Which of the many possible outcomes is most likely to deliver a world of plentiful, high quality, and affordable parts?

Now is surely the time to ask. Academic scientists still control the lion's share of synthetic biology projects, resources, and expertise. Potentially, this gives them important leverage over how industry evolves. But that will change. One company (Amyris Technologies, see below) is already using synthetic biology to make a parts-based organism. Other companies will surely follow, dwarfing the academic sector and eroding its capacity to influence events. At the same time, any action must be carefully planned. Despite similarities to earlier technologies (e.g. classical biotechnology, electronics), synthetic biology is sufficiently unique that interventions based on casual analogies are doomed to fail or, worse, produce unintended consequences. Responsible action must proceed from a deeper and more explicit understanding of how standardized parts will change biotechnology. The natural language for this inquiry is economics.

The tipping dynamic

Compared to electronics, biology research has so far done little to exploit the power of shared, standardized components. However, that is changing. There is already intriguing evidence that companies trying to turn stem cells into, for example, heart or nerve tissue, prefer to start from a small number of standard lines for which they have shared experience (Maurer, 2007). This dynamic is very similar to the electronics industry, where competing companies routinely write programs for a common operating system or game console. We expect such examples to proliferate as biologists' ability to design and manipulate living things expands.

Nowhere is this ambition to standardize more pronounced than in synthetic biology. For now, no one can be sure that the parts agenda will succeed. However, we have already noted that at least one company—Amyris Biotechnology—is betting on success (Ro *et al.*, 2006). According to recent reports (Keasling, 2007), Amyris has already made dramatic progress in making several dozen parts work together inside yeast and bacteria to synthesize the drug precursor artemisinin. Future projects will include rearranging many of the same parts into yeast- or bacteria-based factories for making aromatics and perfumes. These projects should become steadily cheaper as Amyris gains experience by using the same parts over and over again. At least for now, the standard parts idea seems to be paying off.

Given these developments, it is only prudent to think about the economics of an industry in which the parts agenda triumphs. (Of course, the agenda could also fail, but in that case there will be no policy issues to worry about...) Significantly, such an industry must have certain basic characteristics. First, parts will be reused. After all, the parts agenda is pointless if every inventor starts from scratch. Second, assembling standard parts into new organisms will be expensive. Some idea of these costs can be judged from Amyris' artemisinin project. The 5-year, \$20 million experiment reportedly spends 95% of its time trying to find and fix unintended interactions between parts (Keasling, 2007). Third, we expect the cost of using parts to fall each time they are reused, particularly where successive experiments focus on the same metabolic pathways. Based on today's limited evidence, knowledgeable observers think that total project costs could fall by 25% or more the first time parts are reused, and that subsequent reuse will likely cut costs several more times before savings flatten (Keasling, 2006).

In the language of economics, these technology conditions can be neatly summarized by a single statement: ‘The average cost of using a part will decline steeply the more it is used.’ Industries that exhibit this feature have been extensively studied since the 1990s and are now well understood. They exhibit what economists call ‘network effects’: the more users a product has, the more attractive it becomes. Network effects explain the electronics industry’s notorious winner-take-all dynamics and, especially, the rise of dominant firms like Microsoft and eBay.

It is easy to see how similar network effects would arise in synthetic biology: suppose that Parts ‘A’ and ‘B’ perform the same function but that Part A has been used in one more experiment than Part B. Then according to our assumptions, Part A could be 25% cheaper to use than Part B. This makes it much more likely that users will select Part A for future experiments (provided they share in the experiences made by the earlier users of Part A) and generates a ‘rich-get-richer’ dynamic, in which even small differences in initial popularity are rapidly amplified. Indeed, this dynamic works even if Part A is a commercial product and Part B costs nothing. All that is required is for Part A’s owner to set prices lower than Part A’s cost advantage over Part B.

The fact that different parts must be used together also gives actors a strong incentive to create entire libraries of parts, in much the same way that software companies develop multiple programs to cover a range of applications. Suppose that Company Y owns 70% of the most popular parts and Company Z owns the remaining 30%. Then, Company Y can get 100% of the business by offering a complete suite of whatever parts that the users need, that is, turning the contest into a competition between libraries instead of individual parts. This can be done by refusing to license parts individually or by offering volume discounts. In theory, courts might strike such practices down under the antitrust laws (*LePage’s*, 2003). But even if they did, companies would still want to develop competing libraries. As the French economist Augustin Cournot (1838 (NT Bacon, trans. 1897)) first pointed out, two monopolists who sell goods that must be used (if at all) together routinely earn less profit than a single monopolist who sells both products. Knowing this, patent owners would prefer to offer a complete library of parts.

Corporate strategy

If and when Amyris succeeds, new companies are bound to enter the market. How will these entrants view the tipping dynamics described above? Basically, there are two possibilities. First, particularly aggressive and well-placed companies will aspire to become the industry’s parts monopolist. But second, other companies will fear monopolization and do what they can to avoid it.

In the software industry, Linux is an example of how this second group can triumph over the first. It turns out that software modules—like parts—seldom, if ever, work the first time they are assembled. Indeed, testing, debugging, and maintenance reportedly account for fourth-fifths of all software development costs. This means that mass market software companies like Microsoft cannot afford to offer more

than a few permutations of their modules to the public. Conversely, companies that want novel configurations must hire programmers to do the work (Bessen, 2005). But why should these programmers buy high-priced commercial modules at all? In the 1990s, many of these programmers realized that it would be cheaper to write their own modules. For them, working on Linux meant ‘creating code that they will never have to pay someone to use again’ (O’Mahoney, 2003).

Like programmers, synthetic biology companies do not want to split their earnings with a parts monopolist. This should give them a powerful reason to donate workers and resources to a Linux-style ‘open parts’ collaboration. However, the Linux model is only viable in certain circumstances. To the extent that commercial contributors’ reward is limited to their own use of the product, they may decide that developing expensive or little used parts is a bad investment. In this case, relying on ‘own use’ incentives might be a show-stopper for synthetic biology. Synthetic biologists often complain that existing parts are both unsophisticated and in short supply. This observation suggests that own use rewards are insufficient. Ultimately, this is an empirical question. Nevertheless, it is important to recognize that a simple Linux model based on own use incentives might not meet synthetic biology’s parts needs.

Where, then, would the extra reward come from? Absent government or foundation funding, the answer, plainly, is ‘other parts users.’ This, of course, is what patents are supposed to do. However, patents often compound the tendency of network markets to tip into monopoly, technically inferior products and other pitfalls. A much better solution would be to meter protection so that it lasted just long enough to supply the part-maker’s (otherwise missing) reward and then disappeared. While this may sound fanciful, detailed study by one of us (J Henkel) has shown that such systems actually exist. Software firms that write code for ‘embedded Linux’ usually sell their services to a handful of companies that make, say, DVD players or machine tools. However, the Linux General Public License (‘GPL’) does not require them to disclose their code to the general public until devices containing it have reached a mass market. This creates an 18-month window in which the code remains proprietary. Surprisingly, this short protection period seems to provide adequate incentives—if it did not, innovation would have stopped by now. Indeed, the average protection period is even shorter. This is because companies voluntarily reveal about 50% of their code before GPL requires them to (Henkel, 2006). Better still, the business models that drive this behavior are not specific to software and ought to work in synthetic biology (see Box 1).

Making a policy choice

So far we have argued that synthetic biology companies have dollars-and-cents reasons to design parts that are either open or—as in embedded Linux—become open after some short period of time. The fact that something is feasible does not, however, make it desirable. Do we want open source parts at all? Too often, this debate has been framed in terms of vague, *ad hoc* analogies (‘parts are like operating systems...’) or else

Box 1 Open source and the bottom line

Companies in the embedded Linux industry use many business strategies to capture value. Synthetic biology companies can exploit most of them:

Open parts, patented products. Software vendors often share basic parts and modules, while protecting the products made from them. Synthetic biology companies could similarly share parts while patenting completed organisms. Companies that create organisms for individual clients may not need patent protection at all.

Shared development. Software vendors routinely share code, hoping that others will update it, identify and fix bugs, or write extensions. Synthetic biology companies could similarly learn from users and even competitors.

Establishing a user base. Software vendors rely on open source to attract users. Synthetic biology companies similarly want to see their parts used as early and as often as possible.

Other strategies. Software vendors participate in open source to demonstrate technical prowess to would-be clients, learn new product ideas, demonstrate social responsibility, and hire talented programmers. Similar motives should also operate in synthetic biology. Synthetic biology contains almost all of the same ingredients that make embedded Linux successful. First, synthetic biology's parts approach emphasizes strong modularity. This allows the work of creating a parts library to be spread over many companies. It also makes it possible for companies to earn profits by patenting some parts while making others openly available. Second, we expect companies to have fairly idiosyncratic parts needs. This means that they cannot simply 'free ride' by waiting for others to make what they need. It also suggests that companies can often share parts without losing their technological 'edge' to competitors. Third, different companies will have different expertise. This suggests that community-based libraries will often outperform company ones. Finally, the synthetic biology market will probably include large numbers of small, idiosyncratic customers. This makes patent licensing less lucrative and, by comparison, openness more attractive.

Box 2 Are life sciences companies ready for open source?

Economic incentives are relentless. If open source makes money, companies will adopt it sooner or later. That said, progress will be faster if industry is open-minded. Recent events suggest that life sciences companies already understand openness and do not hesitate to use it as a business tactic. Such practices date back to 1999, when 10 pharmaceutical companies funded the SNP Consortium to put genome data in the public domain so that their competitors could not obtain patents to corner the market. [1] More recent examples include Pfizer's decision to disclose the contents of its drug discovery pipeline [2], Syngenta's decision to share its rice genome data [3], and Novartis' decision to release their genome-wide type 2 diabetes map over the Web. [4] So far, this sharing has been largely limited to releases of basic data. At least arguably, parts are closer to a working product, and to that extent would require at least a small leap of imagination. That, however, is what smart businessmen are supposed to do.

[1] Janet Hope, *Open Source Biotechnology* (December 2004) (unpublished PhD dissertation, Australian National University), available at <http://opensource.mit.edu/papers/hope.pdf>.

[2] Anon, 'Pfizer Pipeline as of December 20, 2006' available at http://www.pfizer.com/pfizer/help/download/product_pipeline_view.pdf.

[3] Dennis Normile, 'Syngenta Agrees to Wider Release,' *Science* 296: 1785 (2002).

[4] <http://www.novartis.com/newsroom/news/index.shtml> (item for Feb. 12 2007).

turned on normative goals ('morality of patenting of life forms') that are beyond reason and evidence. Here, we focus on a narrow and well-defined question: What mix of patent and open source incentives is most likely to deliver the cheap and abundant parts that synthetic biologists (and, by implication, the world) needs? (Box 2)

In modern innovation economics, this turns out to be a two-part question. First, 'static efficiency' suggests that, once a part exists, its blueprint should be freely available to everyone. Patents do poorly on this ground because they let inventors charge high prices for inventions. On the other hand, 'dynamic efficiency' suggests that inventors need incentives to create parts in the first place. This often means letting them extract profits from users. Patents handle this problem very naturally (Scotchmer, 2004). These arguments lead to the following observation: if open source incentives can produce a given part, then society is better off without patents. But if not, it is still better for a society to have an expensive, patented part than no part at all. Given our limited knowledge of the world, we cannot write down (or at least cannot prove!) a clear rule, which specifies parts that should be patented and those that should be open. Moreover, even if we could, the patent laws would not allow us implement it.

We can, however, do the next best thing. Writing down an explicit rule may not be necessary if we can create incentives that encourage companies to make the right decision on their own. We have already said that open source is desirable, except where patents are needed to elicit parts production. It follows that society is better off whenever a company creates a part and then voluntarily participates in open development instead of patenting. The trick, of course, is to design institutions that encourage such decisions. The institutions that made embedded Linux work were, in some sense, an accident. Can synthetic biologists foster them intentionally?

Steering toward open source

For the community to exert influence, it must first control something that companies want. Fortunately, it has such an asset. We have stressed that a part's value increases each time it is used. But experience means nothing unless it has been captured and stored. Today, most of the world's parts data is found in the community's Registry of Standard Biological Parts (<http://parts.mit.edu>). Synthetic biology companies will certainly clamor for these data. But what should the rules be?

There are several possibilities (see Rai and Boyle (2007) for issues of legal implementation). The first—and simplest—is to make the data available to anyone for any purpose. Today's Registry does this. Interestingly, even this very minimal intervention makes openness significantly more likely. Recent economics research shows that companies can often earn more money by pooling information instead of hoarding it (Allen, 1983; Nuvolari, 2004). By capturing other users' experience, the Registry would enhance the benefits of pooling.

A second option would be to limit the Registry to open parts. This would reward companies each time they donated a particular part to the public domain. This is basically the embedded Linux model.

Finally, the community could open the Registry to any owner that promised to donate its parts to the public domain after some fixed number of years or, perhaps, royalty payments. This is probably the best option, since companies could always make a business decision to donate their parts early (our second option) if conditions warranted. The main practical problem would be to decide how many years to ask for: too few and partsmakers will stay out of the Registry, too many and the benefits of openness are deferred. One way to avoid these dangers would be to select a conservative number (say, 5 years) and then adjust it downward, based on the number and quality of parts flowing into the Registry.

So far, we have assumed that the Registry will retain its status as the world's premier focal point for recording and sharing parts information. This seems reasonable, since scientific databases—like parts—almost always follow a winner-take-all dynamic in which early frontrunners become larger and more entrenched over time. On the other hand, would-be parts monopolists might try to build their own proprietary databases to challenge the Registry. We think this risk is manageable. For every company that wanted to monopolize parts data, there would be several others trying to block it. The Registry would almost certainly receive their support.

Conclusion

The ongoing debate between open source advocates and patent enthusiasts is entertaining but also paralyzing. For reasons explained above, we cannot say that open source is 'always' better than patents or *vice versa*. Indeed, such statements are almost surely wrong. Suggestions that synthetic

biologists should combine an open 'operating system' layer with a proprietary 'applications' layer are shrewder, but assume that such lines can be drawn in advance. This, too, is a dead end. Discussions that do not lead to action are pointless.

We do not claim that the embedded Linux-type system outlined above would produce an 'ideal' or 'optimal' result. That can almost never be done in economics. What we do claim is that our scheme is feasible, represents an unambiguous improvement compared to either patents or open source in isolation, and is likely to yield substantial results. It is, in short, worth doing.

References

- Allen RC (1983) Collective invention. *J Econ Behav Organ* **4**: 1–24
- Bessen JE (2005) Open Source Software: Free Provision of Complex Public Goods (mimeo). Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=588763
- Cournot AA (1838) *Researches into the Mathematic Principles of the Theory of Wealth*. (Bacon NT, trans. 1897) New York, USA: Macmillan
- Endy D (2005) Foundations for engineering biology. *Nature* **438**: 449–453
- Henkel J (2006) Selective revealing in open innovation processes: the case of embedded Linux. *Res Policy* **35**: 953–969
- Keasling J (2006) Personal communication
- Keasling J (2007) Video lecture on synthetic biology, March 8, 2007. Available at <http://media.coe.berkeley.edu/BIOSECURITY/03082007/BioLC8.asx>
- LePage's Inc. v. Minnesota Mining and Manuf. Co.*, (2003) 324 F.3d 141–182
- Maurer S (2007) Open source biology: finding a niche (or maybe several). *UMKC Law Review* (forthcoming)
- Nuvolari A (2004) Collective invention during the British Industrial Revolution: the case of the Cornish pumping engine. *Cambridge J Econ* **28**: 347–363
- O'Mahoney S (2003) Guarding the commons: how community-managed projects protect their work. *Res Policy* **32**: 1179
- Rai A, Boyle J (2007) Synthetic biology: caught between property rights, the public domain, and the commons. *PLoS Comput Biol* **5**: 389–393
- Ro D-K, Paradise EM, Ouellet M, Fisher KJ, Newman KL, Ndungu JM, Ho KA, Eachus RA, Ham TS, Kirby J, Chang MCY, Withers ST, Shiba Y, Sarpong R, Keasling JD (2006) Production of the antimalarial drug precursor artemisinic acid in engineered yeast. *Nature* **440**: 940–943
- Scotchmer S (2004) *Innovation and Incentives*. Cambridge USA: MIT Press