Comments Welcome
Please do not circulate or quote

The Guild of Designers

Carliss Y. Baldwin
Kim B. Clark

Harvard Business School

This Draft:
February 21, 2002

Our thanks to Masahiko Aoki, Wayne Collier, Sonali Shah, David Sharman, Kevin Sullivan, Jason Woodard, and members of the Negotiations, Organizations and Markets group at Harvard Business School for sharing key insights. We alone are responsible for errors, oversights and faulty reasoning.

Direct correspondence to:
Carliss Y. Baldwin
Harvard Business School
cbaldwin@hbs.edu

1 The Guild Defined

We begin our exploration of institutional forms with something we will refer to as a "guild." The form we describe does not correspond to any particular guild in history. Instead the name is meant to be evocative of voluntary and free associations whose members share some special skill, language or knowledge (and example would be the Spacer Guild in the novel *Dune*).

In this analysis, the defining characteristics of a *guild* are as follows:

- It is an association of designers;
- Designers belong to it by virtue of being able to "speak" and understand the same "language" of design (in a particular domain);
- All members' actions are voluntary: there is no coercion, direction or management of the membership;
- There are no monetary incentives, and no formal contracts within the guild.

Note 1: The boundaries of the guild are unambiguous. Even if guild members are spread throughout the greater economy, we know who is in and who is out. This will be important when it comes time to look at the financial sufficiency and value-added of this institutional form.

Note 2: Membership goes with certain competencies: in that sense it is not voluntary. However, members' actions are voluntary.

Note 3: The "secret" language which only guild members know is a classic example of something that makes actions observable (to guild members) but not verifiable by outside parties.[1] Thus, unless the guild develops its own enforcement mechanisms, members' "contracts" must be self-enforcing, and guild equilibria must be sub-game perfect.

The purpose of the guild is to solve the generic problems of the Q(k) technological substrate (see Institutional Forms, Part 1). Solving these problems requires collective and/or coordinated action.

For now, we will assume that, by virtue of their specialized training, designers have internalized the x-scale of design quality, and have formed rational expectations with respect to R(k) for each of the M known problems. We further assume that R(k) maps into Q(k) as described above. This means that designers' internalized appraisals of the value of solutions correspond *in*

[1] See Grossman and Hart (1986); Hart and Moore (1990); Hart (2001), and citations therein.

rank to what users are actually willing to pay for such solutions.[2] Thus (we assume) generic problems (1) and (2) have been addressed in the course of the designers' training.

It remains for the guild to address problems (3), publishing and transferring solutions; and (4), coordinating designers' choices of target problems.

Suppose the guild offered the following services to all members at a nominal cost:

- a library of solutions to known problems, with appraisals of the x-quality of each published solution;
- real-time problem sign-up sheets, with counters as described below.

Ex ante, that is, before each problem-solving interval, the library serves as the source of the "textbook solution" to each problem. *Ex post*, the library serves as a pooling-and-ranking mechanism for new solutions. The best of the new solutions would be identified as such, and would become the textbook solution for the next round of problem solving.[3]

The sign-up sheets require some further explanation. Basically, we require the guild to provide a public, accessible place or space, where members can indicate their interest in particular problems. Before each problem-solving interval begins, assume the guild posts one sheet for each of the M known problems. Designers would then indicate their intentions of working on a problem by registering on the signup sheet.[4]

2 The Designers' Selection of Problems to be Solved

To solve the basic coordination problem (described in Part 1), the guild does not have to publish the actual names of those who have signed up. But, it does have to publish/broadcast running counts of the total number of people who have signed up to solve each type of problem.

---

[2] **How those expectations emerge, especially as a function of k; how the expectations get to be "rational" or "self-confirming;" and how design quality (the x-scale) and users' willingness to pay (the y-scale) come to be related; these are REALLY deep epistemological questions. I am reserving them for later… or never!**

[3] What happens to inferior solutions? In theory, the library could discard them, or could store them for future detailed analysis. In real design communities, most inferior design solutions are discarded and forgotten, although some are intriguing or puzzling enough to warrant further study. The arbiters of what is kept and what is discarded are usually the designers themselves.

[4] Note parallels to a matching market. Cf. Roth, many articles.

Thus members scanning the signup sheets should be able to see (in real time or with short delays), a vector $<k_{1t}, k_{2t}, \ldots, k_{Mt}>$, where k_{jt} is the number of people who have signed up to solve problem j at time t.

This information, coupled with the designers' knowledge of the payoff function, $R_i(k)$, for each problem, can serve as the basis for a (new) one-step-look-ahead, self-selection rule as follows:

> At the beginning of each problem-solving interval, let each designer "visit" the signup sheets at least once. For simplicity, we assume each arrives at a different time. On arrival, let the designer review the M known problems and the number of people signed up to work on each one. Based on the available data, let the designer "compute" the expected incremental reward from working on each problem:
>
> $$R(k_{1t} + 1) - R(k_{1t}), \ldots , R(k_{Mt} + 1) - R(k_{Mt}) \;;$$
>
> The designer can then select the problem with the maximum expected reward according to this calculation. That problem will be the target problem, if the designer decides to work on any problem during this interval.
>
> Let each designer then compare the utility of working on his or her target problem (labeled "j") with the utility of using the best solution generated by the k_{jt} people who previously signed up. It follows that the designer will voluntarily work on the target problem, if and only if:
>
> $$u(\omega, R(k_{it}), \ldots, R(k_{jt} + 1), \ldots, R(k_{Mt}), e + \Delta e) > u(\omega, R(k_{it}), \ldots, R(k_{jt}), \ldots, R(k_{Mt}), e) \;.$$
>
> If this comparison is positive, the designer will sign up for the i^{th} problem, and its counter will increase by one. Otherwise, the designer will expend less effort and spend the problem-solving interval applying the textbook solution to a randomly selected "problem of the moment." At the end of the problem-solving interval, all designers will obtain the benefit associated with knowing the best new solution to all of the known problems.

Under the assumptions of our model, in the eyes of each designer, every problem has a well-defined integer k^* such that for integers $k > k^*$, his or her expected marginal contribution, $R(k+1) - R(k)$, is not worth the additional effort. Each problem may have a different k^* (cf. Baldwin and Clark, 2000), but, if the designers have identical preferences and perceptions, as we have assumed, they will all agree on the appropriate k^* for each problem.[5]

As long as the counter on the j^{th} signup sheet shows $k_j < k_j^*$, a designer consulting the signup sheets will have an intrinsic motivation to work on that problem. Absent strategic gaming against the signup process (see the analysis of free-ridership below), said designer will sign up for some problem on the list (the one with the highest expected marginal contribution). If, however, for all M problems, $k_j=k_j^*$, then designers consulting the signup sheets will not sign up for any problem: they will expend less effort, and await the publication of solutions.

Given timely information (updating of the counter), this method of matching problems and designers lodged within an institutional form (the guild) can reach an equilibrium wherein:

- each designer will work on zero or one problem;
- all designers' incremental efforts are rewarded in expectation; and
- *at the end of the signup period*, no designer will have incentives to switch from one problem to another, from some problem to no problem, or from no problem to some problem.

The utility derived by the representative designer from the combination of the signup sheets and the library is:

$u(\omega, R_1(k_{1t}^*), \ldots, R_M(k_{Mt}^*), e+\Delta e)$ if the designer works on a problem;

and

$u(\omega, R_1(k_{1t}^*), \ldots, R_M(k_{Mt}^*), e)$ if the designer does not work on a problem.

Notice, however, that in these utilities lurks a free-rider problem: the designers not working on problems get the benefit of the pooled solutions, without expending incremental effort! This problem is endemic to voluntary sign-ups to supply public goods. The strategic interaction can be modeled as a game.

---

[5] Identical perceptions and preferences are not necessary to our argument, but they do simplify the discussion quite a bit!

2.1 Gaming the Signup Sheets

The generic normal form of this game for two players is shown in Figure 1. The two-player game has two Nash equilibria in pure strategies, and one, less-efficient, but more fair, equilibrium in mixed strategies. The mixed strategy equilibrium is also an Evolutionarily Stable Strategy (ESS) (cf. Gintis, p. 152-153, the Cooperative Fishing game).

Figure 1
Two-Player, Public Goods Signup Game with Generic Payoffs

	Designer 2 Solve	Free Ride
Designer 1: Solve	1; 1	1; 2
Free Ride	2; 1	0; 0

Insert explanation. How much do we need to elaborate this? Do we need to go into the designers incentives to participate? Not participating is a strictly dominated strategy.]

Thus there is reason to suspect that the signup-plus-counter mechanism would work, although not always perfectly.

2.2 The Sufficiency of the Guild

Needs to come from voluntary contributions of $ and effort by designers. Costs need to be "low enough."

3 Tournaments within the Guild

The mechanisms outlined in the previous section, which support the pooling of solutions and the coordination of problem-solving choices, rested solely on the designers' motivations to know and to use better solutions. Hence, as envisioned thus far, the library-plus-sign-up services would function anonymously. But a system capable of publishing solutions can easily publish the names of the authors of those solutions. The recognition of authorship coupled with the ranking of solutions in turn introduces *a new motivating force* into the guild system: the designers' desire to win and/or desire for fame in the designers' community.

We assume that some designers are motivated to compete: they will expend both effort and money in order to participate in competition.[6] Any system that ranks outcomes on some measure of quality, and publishes the names of the winners, is an arena for competition. Thus, if a "Guild" publishes and ranks solutions to known problems, every known problem becomes a potential tournament to be won, as well as a problem to be solved.

We can analyze the effects of competition by adding a new element to each designer's utility function. Specifically, let the utility functions have the same dimensions as before, but let the i^{th} designer's expected reward from solving the j^{th} problem be:

$$a_{ij} R_j(k_j + 1) \quad + \quad b_{ij} R_j(k_j + 1)/(k_j + 1) \ ;$$

where:

- $R_j(k)$ is the order-statistic function of the j^{th} problem;
- k_j is the number of *other* designers working on the j^{th} problem (in a given problem solving interval);
- a_{ij} and b_{ij} are the subjective weights the i^{th} designer places on knowing the solution to the j^{th} problem vs. winning the tournament to solve the j^{th} problem; and
- $1/(k_j + 1)$ is the i^{th} designer's (symmetric) probability of winning, if he or she chooses to work on the problem.

The first term in this expression can be thought of as the expected "pleasure" the i^{th} designer will obtain from knowing and using a superior solution to the j^{th} problem. The second term is the expected additional "pleasure" that same designer would obtain if he or she turned out to be the author of the best solution to the j^{th} problem.

The relative weights of knowing vs. winning in the designer's own preferences are captured by the a_{ij} and b_{ij} multipliers. In theory and in fact, these weights will vary from designer to designer and from problem to problem.

---

[6] It is beyond the scope of this investigation to explain why. Some cites would be nice, however!

Up to this point in the analysis, we have been assuming that $a_{ij} = 1$ and $b_{ij} = 0$. In effect, we assumed that all designers have a "pure pleasure of knowing," and no "pleasure of winning." We will now split the class of designers into two types: one which, as before, has a "pure pleasure of knowing"; and another, symmetrically opposite, which has a "pure pleasure of winning": $a_{ij} = 0$ and $b_{ij} = 1$. In what follows, we will call the first type the "Knowers" and the second type the "Competitors." We are interested in how the two types will interact on the technological substrate of Q(k) design processes.[7]

3.1 The "Competitor's" Decision Rule

Our first task is to characterize the task selection and effort allocation rules of the Competitors. First, it is worth noting that unless the Guild publishes the names of the "winning" authors of solutions, the Competitors will not expend incremental effort solving any problem. In the signup game described in Section 2.1, the Competitors will be free-riders. *Thus, the Knowers have incentives to publish the names of winners, in order to create tournaments that elicit effort from the Competitors.*

Suppose the name of the winning designer for each problem is published at the end of the problem-solving interval. For each Competitor, then, an individually rational one-step-look-ahead decision rule is the following:

> Before each problem-solving interval, let the Competitor visit the signup sheets. He or she can then review the M known problems and the number of people signed up to work on each one. Based on the available data, let the Competitor then "compute" the expected incremental reward from working on each problem:
>
> $\ldots R_j(k_j+1)/(k_j+1), \ldots$;

[7] A different, also interesting, way to construct the Tournament Game is to say that all designers have hybrid motives: a_{ij} , $b_{ij} > 0$. Until the Guild publishes the names of authors of the best designs, the competitive motive will be latent, having no outlet. But if b_{ij} is positive for some designers, and if the existence of competition does not reduce a_{ij} for anyone, then (1) all designers will want to have the winners' names published; and (2) those with positive b_{ij} will increase their efforts in response to the tournament structure. In other words, some who would have been free riders in the absence of a tournament, will enter the lists if they can compete to win!

and select the problem having the maximum expected reward according to this calculation. That problem will be the Competitor's target problem, if he or she decides to work on any problem during this interval.

Let each Competitor then compare the utility of working on the target problem with the utility of not working on any problem. He or she will voluntarily work on the target problem, if and only if:

$$u(\omega, \ldots, R_j(k_j+1)/(k_j+1), \ldots, e+\Delta e) > u(\omega, \ldots, 0, \ldots, , e) \quad .$$

If this comparison is positive, the Competitor will sign up to work on the jth problem, and the counter for that problem will increase by one. Otherwise, the he or she will expend less effort, and spend the problem-solving interval applying the textbook solution to a randomly selected "problem of the moment."

Note that in the Competitor's decision-making calculus, winning is everything. The quality of solutions generate by others does not enter his or her utility function. At the end of the problem solving interval, one designer will "win" the tournament in each category, and obtain the benefits of winning. For the Competitors who win, these benefits may include high status or a good reputation among other designers, as well as the "warm glow" of competitive success.

There is quite a lot of humor implicit in the games between Competitors and Knowers, as we hope to show in the next section!

3.2 Tournament Effects on the Allocation of Tasks

Because of well-known tournament effects, if Competitors and Knowers place equal subjective weight on their respective benefits, Competitors will tend to replace Knowers on many of the problem-solving signup sheets. Thus the equilibria in the signup game include a Tom-

Sawyer-like configuration, in which the Competitors work and the Knowers consume their solutions.[8]

To see why this is so, recall that (by assumption) generating a solution to a problem requires the same amount of incremental effort, Δe from a Competitor as from a Knower. Therefore, we can think of there being a subjective "reward hurdle," ρ: if the subjective reward is greater than ρ, a designer in either class will expend incremental effort Δe; if the subjective reward is less than ρ, he or she will not expend the effort.

A Knower will sign up to work on a problem if:

$$R(k+1) - R(k) \geq \rho \quad ;$$

whilst a Competitor will sign up to work if:

$$R(k+1)/(k+1) \geq \rho \quad .$$

We know that the function R(k) is increasing and (quasi)concave in k, and that R(0) = 0. Therefore, for $k \geq 1$:

$$R(k+1) - R(k) < R(k+1)/(k+1) \, . \qquad (1)$$

Note that the inequality is strict.

We can now state the following:

Proposition 1. Given equal subjective weights on their respective benefits (equal ρ):

a. A Competitor will volunteer to solve any problem that a Knower would volunteer to solve; but

b. A Competitor may volunteer to solve some problems that a Knower would abstain from solving.

Proof.

1a. For the j^{th} problem, we have already defined k_j^* for a Knower as the lowest k such that: $R(k_j^*+1) - R(k_j^*) < \rho$. If the value of the counter on the jth signup sheet is k_j^* or greater, a Knower will not sign up to work on that problem. Necessarily:

$$\rho \leq R(k_j^*) - R(k_j^*-1) < R(k_j^*)/k^* \quad .$$

[8] We don't want to rest too heavily on this result as descriptive of the real world: after all we are looking at only two extreme types. But the Tom Sawyer equilibrium does arise in the model world, and it is suggestive of effects that may arise in the real world.

The first inequality follows from the definition of k* as the lowest value of the counter such that a Knower would not volunteer. The second inequality is implied by expression (1) above. According to the Competitor's decision rule, if the j^{th} counter equals k_j-1, and $\rho \leq R(k_j)/k_j$, then the Competitor will volunteer to solve that problem. The series of inequalities above mean that the Competitor's condition is satisfied for all $k_j \leq k_j^*$ - 1. Thus a Competitor will volunteer to solve any problem that a Knower would volunteer to solve.

1b. Define k_j' as the lowest k such that $R(k_j' + 1)/(k_j' + 1) < \rho$. If the value of the counter on the j^{th} problem is k_j' or greater, a Competitor will not sign up to work on that problem. Part **a** of this proposition implies that $k_j^* \leq k_j'$. Both numbers are integers, thus either $k_j^* = k_j'$, or $k_j^* + 1 \leq k_j'$. In the first case, the Knowers and Competitors will stop signing up at the same point; in the second case, Competitors will continue signing up after Knowers have stopped. **QED.**

Now let us divide the known problems into two categories:

- Type A problems have $k_j^* = k_j'$;
- Type B problems have $k_j^* + 1 \leq k_j'$.

The distribution of problems between these two types is an empirical question; however, problems with a high variance of outcomes in their underlying probability distributions are more likely to be Type B.

The next proposition characterizes the problem-solving equilibria for these two types of problems. In it, we assume that:

- Knowers and Competitors arrive at that signups sheets in a random order; and
- Signups are reversible.

Different assumptions as to the microstructure of the signup process will change the equilibria in ways that are straightforward and obvious.

Proposition 2. Given a reversible signup process with random arrivals,

a. Type A problems will be solved by Competitors and Knowers in (approximate) proportion to their presence in the population.

b. Type B problems will be solved preferentially by Competitors.

c. If the population of Competitors is large enough (greater than $\Sigma\, k_j^*$ over all M problems), Type B problems will be solved by Competitors exclusively.

Proof.

2a. With Type A problems, Knowers and Competitors will arrive randomly and signup as long as the counter value is less than k*. At that point, new arrivals from both sets will stop signing up. However, those already committed to solve the problem will not have reason to reverse their individual decisions. The ending value of the counter will be k*.

2b. With Type B problems, random arrivals and signups will proceed until the counter equals k*. At that point, Knowers will stop signing up, but Competitors will continue to do so, causing the counter to exceed k*. Knowers already committed will then want to reverse their decisions, thus driving the counter back toward k*. But their defections will cause newly arrived Competitors to keep signing up.

2c. Let n_K denote the number of Knowers in the class of designers, and n_C the number of Competitors. It is straightforward, if tedious, to show:

1) If $n_C > \sum_{j=1}^{M} k_j'$, then only Competitors will solve Type B problems, and some of them will be idle; Knowers will work on Type A problems only.

2) If $\sum_{j=1}^{M} k_j^* < n_C < \sum_{j=1}^{M} k_j'$, then only Competitors will solve Type B problems, and none of them will be idle; Knowers will work on Type A problems only.

3) If $n_C < \sum_{j=1}^{M} k_j^* < n_C + n_K$, then no Competitors will be idle; Knowers will work on both A and B-type problems; some Knowers will be idle.

4) If $n_C + n_K < \sum_{j=1}^{M} k_j^*$, then Knowers and Competitors will work on both types of problems; and no one in either group will be idle.

The second set of inequalities proves the statement. **QED**

Above, we said that in the absence of tournaments, the Competitors would free-ride on the Knowers. The existence of tournaments upends this relationship: in the presence of tournaments, Proposition 2 shows that it both possible and likely for the Knowers to free-ride on the Competitors!

References

Aoki, Mashiko (1999) "Information and Governance in the Silicon Valley Model," Stanford University, http://wwwecon.stanford.edu/faculty/workp/swp99028.html, viewed Jan 12, 2001.

_______ (2001). *Towards a Comparative Institutional Analysis*, MIT Press, Cambridge, MA.

Baldwin, Carliss Y. and Kim B. Clark (2000). *Design Rules, Volume 1, The Power of Modularity*, MIT Press, Cambridge MA.

Bates, Robert H., Avner Greif, Margaret Levi, Jean-Laurent Rosenthal, Barry Weingast (1998). "Introduction" and "Conclusion," *Analytic Narratives*, Princeton University Press, Princeton, NJ.

Bowles, Samuel, Herbert Gintis and Melissa Osborne (2001). "The Determinants of Earnings: A Behavioral Approach, *Journal of Economic Literature* XXXIX (December): 1137-1176.

Braudel, Fernand (1982). *The Wheels of Commerce*, (tr. Sian Reynolds) Harper & Row, New York, NY.

Dawkins, Richard (1989). *The Selfish Gene, 2nd ed.* Oxford University Press, Oxford UK.

Gintis, Herbert (1976). "The Nature of Labor Exchange and the Theory of Capitalist Production," *Review of Radical Political Economy*, 8(2): 36-54.

________ (2000). *Game Theory Evolving*, Princeton University Press, Princeton, NJ.

Greif, Avner (1994). "Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies," *Journal of Political Economy* 102: 912-50.

________ (1998a). "Historical and Comparative Institutional Analysis," *American Economic Review*, 88 (May): 80-84.

________ (1998b). "Self-Enforcing Political Systems and Economic Growth: Late Medieval Genoa," in *Analytic Narratives*, (ed. Bates *et al*) Princeton University Press, Princeton, NJ.

Lumsden, Charles J. and Edward O. Wilson (1981), *Genes, Mind and Culture*, Harvard University Press, Cambridge, MA.

Maynard Smith, John (1982). *Evolution and the Theory of Games*, Cambridge University Press, Cambridge, UK.

Maynard Smith, J. and G. Price (1973). "The Logic of Animal Conflict," *Nature, London*, 246: 15-18.

Merton, Robert C. (1973) "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4(Spring): 141-183; reprinted in *Continuous Time Finance*, Basil Blackwell, Oxford, UK, 1990.

_________ and Zvi Bodie (1995). "A Conceptual Framework for Analyzing the Financial Environment," in *The Global Financial System: A Functional Perspective*, (ed. Crane *et al*) Harvard Business School Press, Boston, MA.

North, Douglass C. (1990) *Institutions, Institutional Change and Economic Performance*, Cambridge University Press, Cambridge, UK.

Samuelson, Larry (1997) *Evolutionary Games and Equilibrium Selection*, MIT Press, Cambridge, MA.

Samuelson, Paul (1970). "Maximum Principles in Analytical Economics," Nobel Memorial Lecture, December 11, 1970. Reprinted in *Nobel Lectures: Economic Sciences 1969-1980* (ed. Assar Lindbeck) World Scientific, London, UK.

Shapiro, Carl, and Joseph Stiglitz (1984). "Equilibrium Unemployment as a Discipline Device," *American Economic Review*, 74 (3): 433-44.

Williamson, Oliver E. (1985), *The Economic Institutions of Capitalism*, Free Press, New York, NY.