CREDIT SUISSE | FIRST BOSTON

Thought Leader Forum

About Thought Leader Forum | 2003 | Past Forums

Thought Leader Forum > 2003 > Alph Bingham

Search Investment Banking

Go

- Introduction
- Presenters
  - Juan Enriquez
  - Eric Bonabeau
  - Alph Bingham
  - Anita McGahan
  - Steven Strogatz
  - Paul DePodesta
  - Brooke Harrington
  - Dan Schrag
- 2003 Concept Cards

Alph Bingham

Complexity Theory and Pharmaceutical R&D

The pharmaceutical industry has been enormously successful and profitable for over 125 years, but the business model is being challenged. The goal of any company is to increase its "fitness" on the fitness landscape of its industry. We define "fitness" as the probability of survival. Each company wants to climb to the peak of the landscape, but no one knows where the peaks really are. The map does not pre-exist, and the map is dynamic—it is constantly changing as different players (competitors, partners, vendors and customers) co-evolve and adapt on the landscape.

Most of the Eli Lilly organization is charged with increasing productivity. This helps the organization climb up the local peak on the fitness landscape. This strategy does not necessarily help you find the highest peak—if you are climbing upwards in the Himalayas, but you are not on the slopes of Mount Everest, you will never find the highest peak.

Our group, e.Lilly, was created to explore the radical changes necessary to find other, higher fitness peaks. This raises the question that Michael Mauboussin asked this morning: How much alignment can we afford, but how much misalignment can we tolerate? Getting this balance right is very difficult but very important. Part of the organization is the explorers and other parts are the exploiters. My group is the explorer part of the organization.

After some conversations with Stuart Kauffman in Santa Fe, we identified seven rules that would sound great to most managers but would ensure our failure as explorers. If we forced everyone in the organization to keep the global corporate objectives in mind, for example, we would never be able to venture off of our current local peak on the fitness landscape.

Related Links

- Biography

Concept Cards

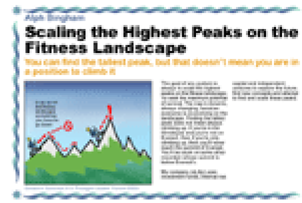

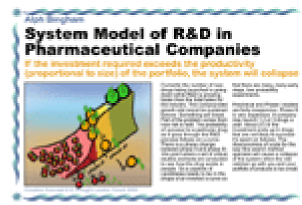

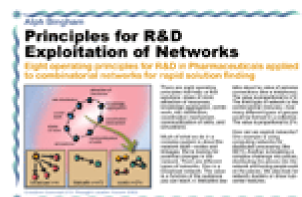

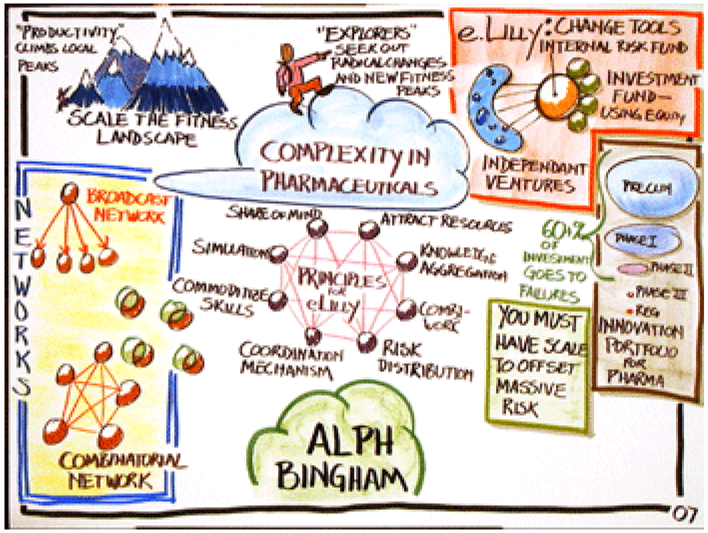

So, what does e.Lilly do? I will admit that the name of the company has worn pretty thin since most e-companies have taken a turn for the worse. Eli Lilly was the name of our founder, and we thought that using his initial as the name of our company was pretty cute. We are charged with transforming Eli Lilly, and I have focused primarily on transforming the R&D organization.

We use several strategies to explore the fitness landscape. The first is an investment fund strategy. We have seven or eight portfolio companies. We try to get a seat on the Board so that we can watch what is happening in the companies from the front lines. The second strategy is our internal risk capital—these are funds that we take to other divisions of Lilly to pursue radical or unlikely ideas that they have developed. Finally, we also create, incubate and spin out new companies, ideas and business models that do not yet exist. These are our three mechanisms for exploring the future, even though most of Eli Lilly is focused on exploiting the current business model and economic structure.

The current operating model in our industry shows an impending challenge. The R&D costs for the industry are growing faster than its revenues. The compounded annual growth rate for R&D between 1980 and 2000 is 14.1% while the compounded annual growth rate for sales over the same period is 9.1%. In the early 1980's we spent $0.10 of every sales dollar on R&D. Today we spend $0.20 of every sales dollar. The right answer might be $0.30—we don't know. We do know, however, that this compounded growth rate cannot be sustained ad infinitum. Something is going to break.

A lot of the problem has to do with the way risk is held within individual pharmaceutical companies. I will talk about how risk is held today, and then I will talk about how network theories and complex adaptive systems may foreshadow a future state very different than the one today.

The probability of the success of a drug is different at different stages of the development process. In the Preclinical and Phase I stages, a drug has a lower probability of making it to the next stage. The probability of a drug's success suddenly skyrockets between Phases II and III. A drug undergoes a set of rigorous clinical trials, and if it proves that it can succeed in humans, then it might have an 80% chance of becoming an approved drug. Many molecules are eliminated at this stage of the process, but the ones that survive have a much higher probability of getting to market.

If you have a pipeline or portfolio of candidate molecules, you therefore must have a huge population of candidates in the early stages of the R&D process and a very few candidates making the final march to the market. For a large pharmaceutical company, their portfolio includes a large number of early-stage, low-probability-of-success compounds, and a handful of compounds going through Phase III and Regulatory.

It does not cost a lot per candidate to do a Preclinical study, and it is very expensive to do a Phase III study. Early in the development process, the R&D costs consist of small investments spread across many different products. Later in the process, all of your costs are focused on the handful of eggs in your basket—the few compounds that you are going to take to market. A company can only launch 1-3 new products per year, so this later-stage portfolio is very small.

If you look at total investment for each stage, then, you discover that about 2/3 of your R&D expenditures go towards compounds that will never make it to market. The majority of the pharmaceutical industry's investment in R&D goes toward failures, but in this business model, this is what you need to invest to ensure that you have successes coming out the other end of your pipeline.

This is a business where most of the risk is held internally. The traditional strategies to manage this risk include diversification and maintaining a large portfolio. You need to test a lot of compounds early so that you may have a handful of products go to market. Unfortunately, you eventually get to the size where this process becomes unmanageable. There are dis-economies of scale for operating R&D this way that make it impossible to get enough scale to diversify the risk across the portfolio. This is why companies like Upjohn, Pharmacia, and Warner Lambert no longer exist—these are companies whose portfolios consisted of maybe fifty compounds. Their risk caught up with them.

At e.Lilly, we constructed eight operating principles to guide our exploration. We need to use language that would be familiar to our culture, but the principles are based on rules for managing complex adaptive systems. To manage a portfolio of exploratory projects, we connect each venture with a handful of these operating principles. Each venture allows us to explore a small set of our operating principles.

Much of what we are focused on is the network itself. I am using "network" in its most general sense—a series of nodes with linkages between them. These links could be social interactions or feedback mechanisms. We want to understand that collective of nodes and relationships between them, and we want to imagine how these kinds of networks might inform us about how we might change the network that we have today. It would be unfair to call a major corporation today anything other than a network.

There are several different kinds of networks. Sarnoff's Law talks about broadcast networks—one node communicating outward to the other nodes. The value or efficiency of this network is roughly a function of the size of the audience you can reach. Mass marketing and advertising are familiar examples of this kind of network.

Metcalf's Law talks about the paired relationships in a network. A familiar example is the telephone system, where individual nodes can pair up to communicate. The value of this network is approximately the square of the number of nodes. This network is of much higher value than a broadcast network.

A third kind of network involved all of the combinatorial mixtures among nodes in a network. How many different groups can you form in a network, from individual nodes, to pairs, to groups of five or ten, to the entire network as a group? The value of this network is even higher still, approximately 2^n. It is in this kind of network that we want to understand how to create value. This kind of network gives us the opportunity to generate enormous increases in value over today's model.

There are already experiments going on in leveraging networks. The SETI experiment is a familiar example. The Search for Extra Terrestrial Intelligence is a grid-computing model that uses unused cycles on other people's computers to analyze radio signals looking for patterns. The search for smallpox treatments uses the same approach to model complex protein-protein interactions. These models are so complex that they would wear down any individual computer. Instead they break the model down into individual problems and distribute those smaller problems into the network of individuals' computers.

We have developed a relatively simple model by applying the learning from these emerging fields to our business. When we launch a new drug product, we naturally try to enroll the opinion leaders around the world. We give them an experience of the drug. The theory is that if the drug is truly safe and these leaders experience that, then they will pass along this information and change the prescribing habits of other physicians. This has been our assumption for many years, but we do not know for sure that this is how the system really works. We therefore applied the lessons of networks to the prescribing community in Boston. We mapped the real influences on a doctor's prescribing practices.

We mapped this network for a drug that treats sepsis. It is an expensive drug created by recombinant DNA methodology, and there are some post-translational modifications that must be made to the molecule that require us to do it in mammalian cells. This drug requires one of the most complex production processes imaginable, and it is consequently an expensive drug. It is intended to be used in critical victims in an ICU with a 35% chance of mortality. One out of three of these people do not go home. Our drug is able to reduce the mortality rate in this population by about 20%, which means that many of them still do not go home, but some do.

Those are the facts about this drug, but how does a physician know whether the drug would help a particular patient, or if they would already be part of the surviving population? How does a physician make this prescribing decision in an environment like this? How is a physician influenced to make this decision?

As we constructed this map, we learned that global opinion leaders were not controlling 75% of local prescribing decisions. Local experiences were far more important. Physicians would build up local knowledge and experiences of a drug, and they would share those experiences with one another through social networks. Our map shows that there were a very few major hubs of influence in Boston. It turns out that one of these hubs was not even on our call list—they were not even receiving our marketing information. We had completely missed one of the major sources of decision-making power in this community. This is a case in which simply creating the map created changes in the behaviors of the Eli Lilly organization.

The next example is a bit more complicated. We created a different map, and then tried to inject some very simple rules of behavior. It is very difficult to determine what these kinds of rules should be, so we used the very simple trading rule, "Buy low. Sell high." With nothing more than this simple rule, we were able to explore how markets could aggregate information for us. This is the model that the Pentagon tried to use to aggregate terrorism information.

Our model of markets is less politically sensitive than the terrorism market, but it is not without its risks and challenges. What if we set prices in a stock market by allowing the CEO to set his company's price every morning? Or what if we allowed three of the big brokerage houses to collectively set prices each morning? This would not appear to be a good way to make decisions, would it? We prefer to have many free agents working with a simple rule and allow the price of a company to evolve.

Inside of most organizations, however, decision makers actually like the idea of making all of their own decisions unilaterally. Leaders might allow decision-making to evolve into a small committee structure, but that's about as far as it goes. The prospect of allowing decisions to be made in an open market format is terrifying to most corporate executives. Lone decision-makers can always rationalize and explain their decisions. Executives are far less comfortable explaining a decision by saying that the market "suggested" one approach over another.

Artificial markets create allocation efficiency and information efficiency. Allocation efficiency is the flow of goods and services from the lowest cost to the highest value users. eBay is a good example of this. The value of a 1937 Radio Flyer wagon used to be set by either an antique dealer in Manhattan ($500) or a farmer in Missouri holding a garage sale ($0.50). By creating a market, eBay has created allocation efficiency.

Information efficiency is about how the market aggregates the knowledge and beliefs of all of the traders in that market. Some information markets will play a role in knowledge management. My concern about knowledge management systems to date is that most of their structures do not seem very knowledgeable. They are heavy on memory and low on cognition. The gather and organize a lot of information, but there seems to be little cognitive process or thinking. This component is required if these systems are to be called "knowledge" management systems. The next example shows some glimmers of promise for a cognitive system that could be laid on top of today's knowledge management systems.

Suppose I created a share of stock that will pay $100 if there is a military coup in Argentina before 2006. It will pay $0 if 2006 dawns and there is no military coup in Argentina. Who will give me $99.50 for that share of stock? No one? Ok, who will give me $0.50? There are lots of takers! There is very good upside potential and very little risk.

Now, for those of you who bought the stock, who would sell it for $1? How about $5? This group is not particularly knowledgeable about South American politics, but you all have bits of knowledge. During this activity, you began to aggregate those bits of knowledge. We may have ended up with a share price of around $6.50 or so. It would not have been over $50 per share, and it would not have dropped under $1. This begins to reflect the rough probability of a military coup in Argentina. This is the nature of the market that the Pentagon tried to establish for terrorism.

This kind of market has been tested in a variety of instances already. A market was established to predict the 1988 election, and the share price was created to match the bias in the popular vote. The final result was far closer to the final result than any of the public opinion polls taken on the night of the election, even though this is when these polls should have been at their maximum accuracy.

In our experiment, we used six hypothetical drug candidates in order to shake down our processes and trading mechanisms. We turned all of our project management staff and a few other members of the organization into our traders. We then took all kinds of data about these six hypothetical drugs and fed it out to our traders in different ways. Not everyone got all of the same information. We knew in advance which drug candidates would survive, and we want to test to see how quickly our market would aggregate this data to come to the same conclusions. The candidates divided into two populations well before the final test was completed on the drugs. The three high-trading candidates were indeed the three successful drugs. The three low-trading candidates were the failures. Because we were also able to monitor what each trader was willing to pay for each drug, we were able to also monitor the "consensus" opinion of the organization about each of the drug candidates—this would be the result of averaging the opinions of the entire organization about each drug. It takes a very long time for consensus opinion to catch up to the prediction of the trading price.

In real practice, this shows a very interesting dynamic within organizations. If you ask a scientist what his opinion of drug X is, he will tell you what a promising and intriguing concept it is, and how hopeful he feels about its success. You will end up with a very positive impression of the drug. If he holds shares in the drug, he'll sell them. This gives you a much better idea of how he values the drug. Taken across an entire organization, this may be the most valuable information that you can acquire.

We have created a map, and we have layered on a few simple rules to extract certain behaviors. The next step is to create an entirely new business around these network theories. The first principle that we wanted to leverage is "Share of Mind". Share of Mind has to do with where the phase changes occur in a complex system as the size of the network enlarges. There are mechanisms here that allow us to perform complex searches and complex evaluations. The best example we all know is about is talk radio. In Indianapolis in May, there is a talk radio program about the Indy 500. A person calls in and asks who had placed third in the race in 1938. This is a very obscure question yet someone else calls in with the answer almost instantaneously. This is extremely weird behavior, and businesses almost never tap into this kind of behavior. If there is a sufficiently large population, then I can find a near-instantaneous answer to an obscure or difficult question. I only need to find a way to reach a critical mass of people for this behavior to emerge.

You are starting to see this behavior in help lines. Most businesses have not tapped into these help lines. I had a problem with my Nikon D100 camera and lens compatibility. I went through many frustrating exchanges with Nikon, and their final response was for me to send them my entire camera and peripherals. I decided to go to a users community for help instead. This community includes more people than Nikon could ever hope to hire. I sent my problem into this community, and the first response came back in 22 hours and gave me a viable solution to my problem. I sent him a little bit more information, and the second response came back in 30 minutes, and gave me the complete solution to my problem. Nikon cannot afford to hire enough people to create this kind of critical mass of people and knowledge.

The second principle that we wanted to explore was Risk Distribution. One of the big challenges that the pharmaceutical industry faces is the way in which risk is held internally in these companies. This has been the downfall of many pharmaceutical companies.

Risk has several important qualities. First, when risk is spread out thin enough and cheap enough, then the upside is always more attractive than the minimal downside. You were all willing to pay fifty cents for that Argentinean deal, even though you didn't know anything about Argentina. You are all money managers, yet most of you have probably bought a lottery ticket before. You know the odds of winning, and yet you still played.

You would probably not have bought the ticket if the price of the ticket had been $1,000. The risk is broken down so far in a lottery that it is not an expensive risk you are taking. The second factor is that a lottery ticket is worth more than a dollar. In addition to the net present value of a lottery ticket, which might be $0.13, you derive $0.95 of daydreaming. This brings the total value of the ticket to $1.08. You only paid a dollar, so your return is 8%. Therefore a smart money manager will buy a lottery ticket.

You must therefore look at all of the utility functions within your network. When you look at a large enough network, you can adopt utility functions that are not available to you in a structured environment such as a corporation. Because there are so many other utility functions in play, you discover that the risk is actually asymmetric—the risk that I offload is not necessarily equal to the risk that you take on. When you look at our industry, you see that the majority of the risk in held internally and that the majority of the R&D funds go to failure. You might ask why our network would have any interest in paying for our risk. Our network, however, does not necessarily assume the same risk. They might be a professor of biochemistry whose mother died of a brain tumor and who has a personal passion for curing this disease.

Innocentive is the company that we have created to explore these network principles. Innocentive runs a website that seeks out solutions to scientific problems. We posted a challenge to create a chemical synthesis for a complex organic molecule. Lilly has 300-600 chemists internally. If we really worked hard at it, we might get 300 people to look at this challenge. Practically speaking, we normally get about six minds to work on any particular drug problem that we have. This challenge involved a very large search space. Libraria has created a database to categorize these kinds of reactions, and they have listed about 14,000 classes of such complex reactions to date. If it takes about ten different reactions to get to a final molecule, then our search space has expanded to 14,00010. Not all of those are viable, so we can eliminate some. But then there are probably 1060 different combinations of starting materials. The combinatorial space is exploding for this problem. It is very unlikely that we will be able to search this space in a programmed and algorithmic way.

Now, imagine the following hypothetical situation. You have an infinite number of graduate students taking a course in organic synthesis. They are given a problem. Let's say that you might get several million potential answers. Most of them would receive a D or F. Several thousand would receive a B, and several hundred would receive a high mark in the course. This is still an enormous search space, especially if you consider that maybe six people in our company are likely to look for an answer.

Instead, post your challenge in the Innocentive environment. We have 28,000 scientists around the world. This is almost two orders of magnitude more scientists focused on this problem than we can even afford to hire into our company. These scientists perform the same search of the same space, completely at their own risk. We had no idea whether these people would perform this search or sort of results we might get, so we launched the company. It turns out that scientists will indeed do this work at their own risk.

We launched the site in June 2001. We all know the events of September 2001, and by December 2001, we were receiving all kinds of packages containing white powders from previously unknown addresses. We were delighted! This was the best thing we could imagine. We started to push the boundaries. Would they send 50 grams of a particular powder? It was a complex molecular structure that is not commercially available. We offered $50,000 for 50 grams, and it showed up in the mail. There are a number of IP and security issues, but we deal with this in various ways. We publish in Chinese, Spanish, English, Russian and German, primarily because these are the dominant languages for science.

The traditional pools of information in this kind of case has been the Nobel Laureates, but among the hundreds of thousands of other scientists out there, could there not be a person better qualified to answer a particular question at a particular moment? This is our way of looking for the "happily prepared mind". Better yet, this is a way to help them find us.

There is an enormous set of reasons why a scientist would play in a network like this one. Take the former chief scientist of Hoechst Celanese. This was a guy who loved working in the lab, but he kept getting promoted instead. When he retired, he used his stock options to buy a laboratory. Instead of playing golf, he putters around the laboratory for about four hours every morning. This is his way of continuing to feel vital and alive. Our website provides him with his problem sets. Our challenges make him feel purposeful. He was one of our first awardees.

These scientists come from all over the world. We once received a solution in 72 hours from Kazakhstan. This is another quality of this network mechanism—sometimes the problems have already been solved. These scientists hadn't solved the problem in 72 hours. They already had it solved, probably for some other reason.

The solver community is global. The percentage of this population that is in the US is shrinking quickly, from 61% a year ago to 47% today. Because of our communications infrastructure, this concept initially exploded in the US. After that initial explosion, however, other countries have a very different vesting in this concept. One of the ways that risk is distributed is that this is a bounty-hunting model. In our approach, we offer a cash reward for the winner, but we assume no responsibility for people who tried and did not win.

This is our deliberate attempt to offload some of our risk, but we have also set up a business model where we must manage the risk on behalf of the risk-taker. The bounty is the global bounty—you get paid the same prize no matter where you live. A scientist in a small village in India solved a $75,000 problem. He is now the local rock star-no one could have imagined than amount of cash coming into that economy. A single, global price helps to manage the risk for our network. The second risk-management strategy is that the intermediary company, Innocentive, does not get paid unless the solver gets paid. This incents Innocentive to protect the solvers from IP theft and piracy, and to nurture a community of solvers. This is a network of 28,000 non-employees that we must protect and nurture. The global bounty helps to explain why our network is growing so quickly in places like India and Russia. The other lure for scientists is the global access to problem sets. These people can get access to companies' problems that would otherwise be impossible to approach. The US is still in the lead, followed by India, China and then Russia. These are communities of strong scientific excellence, but they do not necessarily have strong access to Western problems or rewards.

Companies that are posting problems to this site tend to be large, household name corporations. A few of them prefer to remain anonymous.

The stories of our problem-solvers are all fascinating. The Indian rock star is a favorite. The former chief scientist has been a solver for us. One solver had been unemployed and managed to solve a problem outside of his field. Other solvers are people in research organizations who have an ebb and flow of internal projects—they use our projects to fill in the "unused cycles" in their business pipeline. These are the people that we thought would be antagonized by our approach, and it has turned out to be a boon for them.

I have not talked about the businesses that we are still incubating. We are trying to link this concept of risk distribution with market mechanisms like securitization. This would allow us to distribute not only the risk of solving a particular problem, but also the risk that that problem would ever become a product by securitizing future royalties as a way to pay for our current R&D activities. We are taking the ideas that you have seen and layering on other principles that I have not discussed in much detail. We are exploring these ideas to discover what they might mean to Eli

Lilly's underlying business.

One source of inspiration for those few of us in this exploratory organization (and a source of antagonization for those people in the exploiting part of the organization) is this quote from Peter Drucker: "The corporation as we know it is now 125 years old, but is unlikely to survive the next 25 years." Legally and financially, there will still be a structure. But organizationally and economically, these creatures will look much different. We are trying to get a glimpse through these experiments of what the future organization might look like.