

When development costs are high and failure is common, companies should structure research to seek truth first, success second. Experiments at Eli Lilly show how it's done.

A More Rational Approach to New-Product Development

by Eric Bonabeau,
Neil Bodick, and
Robert W. Armstrong

COMPANIES OFTEN TREAT new-product development as a monolithic process, but it can be more rationally divided into two distinct stages: a truth-seeking early stage, focused on evaluating novel products' prospects and eliminating bad bets, and a success-seeking late stage, focused on maximizing the value of products that have been cleared for development. Recognizing the potential of this approach, in 2001 Eli Lilly designed and piloted Chorus, an autonomous experimental unit dedicated solely to early-stage drug development. Chorus looks for the most likely winners in a portfolio of molecules (most of which are destined to fail), recommending only the strongest candidates for costly late-stage development.

The unit has evaluated 19 drug candidates, 12 of which are still being assessed. (By the end of 2007, Chorus had completed work on seven molecules, recommending that four enter into full-scale clinical development and that the other three go no further.) Although Chorus absorbs just one-tenth of Lilly's investment in early-stage development, it has recently delivered a substantially greater fraction of the molecules slated for late Phase II trials - at almost twice the speed and less than a third of the cost of the standard process, in some cases shaving 12 to 24 months off the usual development time.

The success of Chorus represents the ideal match of an innovation-management problem and solution. The model is well suited to drug development because, although it may postpone the scale-up of successful products, it reduces risk in an environment where development costs and failure rates are extremely high. Indeed, any company that needs to absorb a lot of risk in early-stage development—for instance, in the chemical, biotechnology, medical devices, high-technology, and semiconductor industries—could probably benefit from adopting the Chorus model. The model would make less sense for companies that have low development

costs and failure rates and are therefore well served by concurrent engineering or rapid-prototyping approaches that promote fast scale-up at relatively low risk.

Consider, for example, how two different molecules were evaluated in early development. In 2001, Lilly had begun work on a drug candidate for treating psychosis that we'll call molecule X32. Three years later, human brain-imaging studies showed that little of the drug actually reached the central nervous system—in all likelihood, not enough to have a therapeutic effect. Nonetheless, the development team kept the project alive, arguing that only minute amounts of the molecule should be necessary to get results.

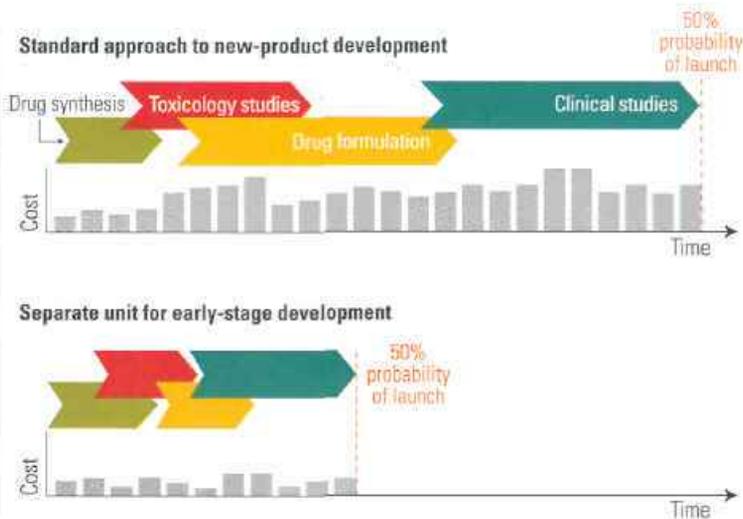
Fast-forward to 2006. After five years of conventional development, it was still unclear whether X32 had any clinical promise. Frustrated by the lack of definitive information, Lilly managers handed the molecule over to Chorus for evaluation. Chorus undertook a new set of small-scale clinical experiments and in just seven months demonstrated that X32 had no therapeutic benefit. This put an end to years of costly procrastination. The resolution was quick, decisive, and obviously cost-effective.

Determining Probability of Launch

Here are alternative plans for developing a single pain-relief drug candidate. The upper chart represents typical late-stage-oriented, success-seeking behavior; the lower one shows an early-stage, truth-seeking approach—the kind Chorus employs at Eli Lilly. Companies with a separate organization for early-stage development can determine more quickly and less expensively—and with a similar level of probability—whether a product will launch.

The chevrons represent segments of work, and the gray bars indicate the associated costs. The effect of ongoing work in each plan is to either increase or decrease the probability of launch. Typically, the clinical test of safety and efficacy (the chevron on the far right) has the greatest impact on launch probability.

In a success-seeking program, expensive and lengthy large-scale manufacturing and long-term animal studies are often initiated before critical data from the early-stage safety and efficacy studies are available. So in this hypothetical case, in the standard success-seeking path an extensive effort (yellow chevron) is made to manufacture a sustained-release tablet to clinically test effects on pain—the crucial



experiment. In the truth-seeking plan, minimal work is conducted to support the crucial experiment. Instead of being delivered as a manufactured sustained-release tablet, the drug is repeatedly administered as a suspension in de-gassed Sprite, which mimics the effects of sustained release in a fraction of the time and at much lower cost. In addition, measuring the drug's effect on a surrogate marker reduces both the cost of the crucial experiment and the time needed.

Meanwhile, Lilly managers turned to Chorus to reevaluate a second drug - 4AB, for short - that had looked promising for certain neurological disorders but had been abandoned prior to clinical testing because similar molecules were found to affect vision at therapeutic doses. Tapping a network of in-house scientists and external academics, Chorus identified a novel biomarker to help in testing the compound's efficacy. The unit then ran several small trials, finding that 4AB did not cause visual problems and was likely to be of clinical benefit. Chorus's new data put 4AB back in the running, motivating large-scale investment in further clinical testing. The drug is now in late Phase 11 trials, and preliminary data suggest that it is both safe and effective.

Chorus delivered these results by focusing on what should be the only objective of early-stage development: reducing uncertainty about a drug candidate's clinical promise - or lack thereof - quickly and effectively.

Kill or Persist?

The examples of X32 and 4AB illustrate two classes of decision-making errors that can impede traditional drug development and new-product development (NPD) in general. One type occurs when managers ignore evidence challenging their assumption that a project will succeed. There are many reasons for this sort of failure, including the power of champions to stir up collective faith in a project's promise and the human tendency to seek only evidence that supports our beliefs. Projects like X32 that survive despite multiple red flags are the outcome; some of them even reach the market, only to fail dramatically after their introduction.

The other type of error occurs when a project is terminated prematurely for lack of evidence that it could succeed. Such mistakes result from a failure to conduct the right experiments to reveal a product's potential, sometimes because of organizational or personal biases against the project or because of a shortage of resources. Halting the development of 4AB falls into this category. Indeed, some of the pharmaceutical industry's biggest blockbusters, such as Prozac, narrowly escaped cancellation due to this kind of error.

Neither class of error is unique to pharmaceutical development. The first type, ignored evidence, abounds in industries ranging from chemicals to building materials to entertainment, where new products with questionable viability - remember RCA's videodisk? - are propelled to market by a dogmatic, success-seeking mentality. (For more examples, see Isabelle Royer, "Why Bad Projects Are So Hard to Kill," HBR February 2003.) And many mature companies cancel promising projects too early for lack of adequate data. Xerox,

Eric Bonabeau (enc@icosystem.com) is the chairman of Icosystem, a strategy-consulting firm in Cambridge, Massachusetts. Neil Bodick (nbodick@gmail.com), recently retired, was the chief operating officer of Chorus, an R&D unit devoted to early-stage development at Lilly Research Laboratories in Indianapolis. Robert W. Armstrong (roba@lilly.com) is the vice president of global external R&D for Lilly Research Laboratories.



for example, abandoned projects that went on to drive the success of Documentum and 3Com.

Any company in an industry that relies on NPD for growth must avoid both kinds of errors. This requires encouraging what may seem like contradictory instincts: a willingness to kill a product early and a willingness to persist until its potential is realized. Management consultants and portfolio theorists have offered a range of opinions on the shortcomings of NPD in large organizations, but none have managed to address how to avoid both types of decision-making errors simultaneously.

That's because most organizations promote both kinds of errors by focusing disproportionately on late-stage development; they lack the early, truth-seeking functions

whose explicit job is to head off such errors. The late-stage model - which in drug development is designed for massive pre- and postlaunch activities - imposes a rigid bureaucracy that encourages large-scale experiments, conducted to maximize the likelihood of launch. For many large companies, this approach comes naturally, because their NPD objectives, incentives, processes, and workflows are geared toward seeking success. But this makes it hard to expose the truth about risky prospects quickly and cost-effectively. Because a late-stage mind-set dominates most innovation companies, creating an early-stage organization with its own objectives, governance, and operations often requires a fundamentally new way of thinking.

Building an Early-Stage Organization

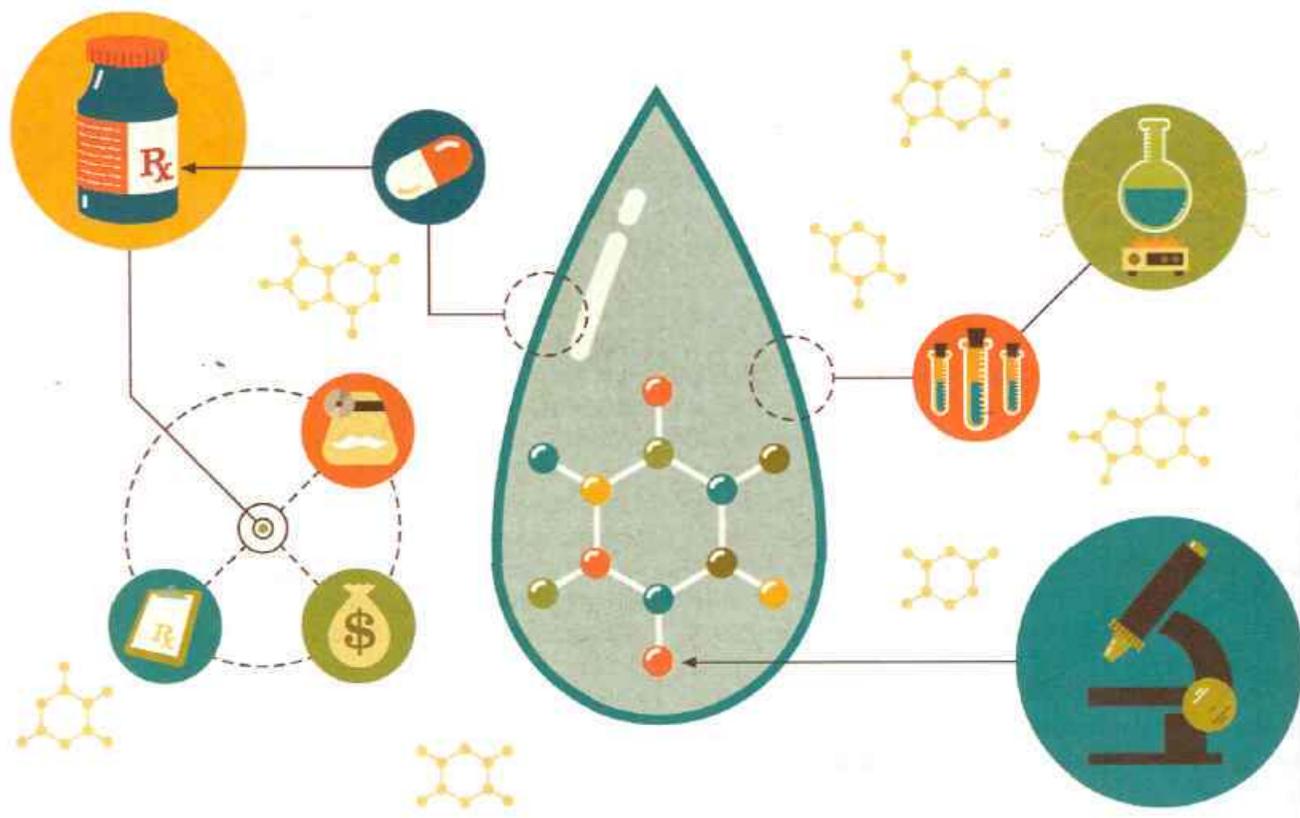
Chorus defines "early stage" as the work of determining proof of concept (POC) for a drug candidate. Researchers must show - in small, highly focused clinical trials - that the drug is likely to be effective and not to have obvious serious side effects. Establishing POC reduces uncertainty about the product's prospects for commercialization and measurably affects the probability of launch.

Unlike the late-stage organization's portfolio, which consists of products headed toward launch, Chorus's portfolio is made up of experiments conducted primarily to resolve uncertainty about a drug candidate's promise and thus substan-

tially increase or decrease the probability that the candidate will launch (see the exhibit "The Two Faces of Pharmaceutical New-Product Development"). Changing this probability involves first identifying key attributes that would affect commercialization (for example, Does the drug occupy and affect its biological target? Does it show efficacy? Does it have undesirable side effects?) and then designing small experiments to establish whether these attributes exist. As data flow from the experiments, Chorus managers modify the experimental plan weekly or even daily in order to discover the intrinsic attributes of a candidate as efficiently as possible.

Because experiments are valued according to their impact in determining the probability of launch, whether they increase it is immaterial to Chorus. The staff cultivates loyalty to the experiment, not to the product. Failure, then, is not only acceptable but periodically expected and rewarded. Reducing uncertainty quickly and inexpensively is the goal that drives the Chorus process, which consists of defining what data are required to change the probability of success, designing the simplest clinical trials that will provide such data, executing the trials cost-effectively, evaluating the data objectively, and delivering a recommendation to either continue or terminate development.

Although Chorus's approach is novel, the notion of pursuing such high-yield "killer," or critical, experiments is not new. About 14 years ago, P. Roy Vagelos, then the CEO of



Merck, lamented the fear that such trials inspire. In an interview in *Harvard Business Review* (November-December 1994), Vagelos observed: "There is one sure road to failure that I have seen many wander down: some people become so afraid of failing that they are unable to do a critical experiment....[Merck] has missed out on some major opportunities because people were unwilling to take that truth-telling step-to conduct the experiment that would show once and for all if what they had spent so many years studying would actually produce a new drug."

Efficiency also requires avoiding large fluctuations in resource utilization, the bane of new-product development in general and early pharmaceutical development in particular. To prevent idle capacity, Chorus taps a network of 50 external experts, who advise on topics such as experimental design and drug delivery, and 75 external vendors, who provide most of the manufacturing, toxicology, and clinical work the unit requires. This frees Chorus's staff of 24 (15 of whom are senior scientists) to focus on the evidence generated by the trials. As a result, 80% of Chorus's annual expenditures are dispersed through the network; the remaining 20% are the fixed costs of running the unit. In addition to providing flexible capacity, such outsourcing reinforces truth-seeking by injecting dispassionate outside perspectives.

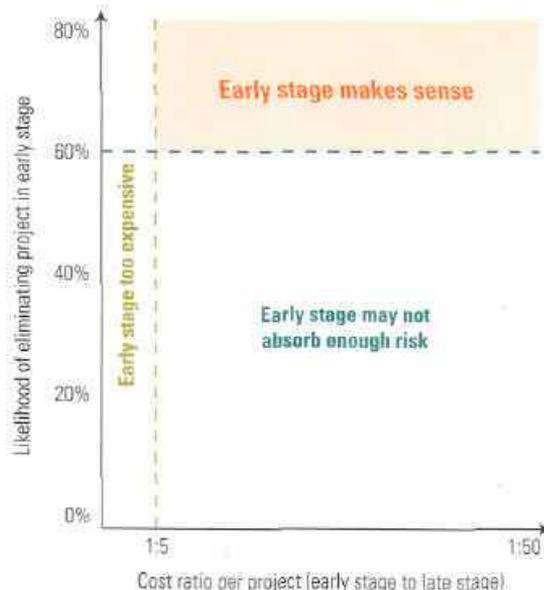
The considerably complex job of managing the work of vendors and outside experts with minimal in-house staff is facilitated by a suite of software tools developed for the Chorus enterprise. At the level of the portfolio, the software suite, known as Voice, tracks the impact of different experiments on probability of launch; at the level of planning, it integrates the opinions of external content experts; and at the level of operations, it organizes work according to subject area (clinical, toxicology, manufacturing, and so on) and distributes tasks and associated documents throughout the network.

A Choice of Models

The Chorus model can help companies improve the efficiency of their innovation processes by establishing proof of concept early and reducing project attrition downstream, particularly in the later and more expensive phases of drug development. However, such truth seeking does have a cost: It may impede parallel processing or concurrent engineering and defer scale-up and commercialization of products that will ultimately prove successful. For example, in a Chorus experiment it is possible to use a test molecule made through an unoptimized process that would not be adequate for larger-scale trials and commercialization, but waiting until Chorus delivers a POC before starting the time- and resource-intensive optimization could delay launch and hinder commercial success. Nonetheless, the net benefit may be substantial. In large pharmaceutical firms, 80% to 90% of drug candidates that enter clinical trials will never launch;

When a Separate Early Stage Makes Sense

You may be able to increase your chances of success in new-product development by dividing the process into early and late stages. In the former, the goal is to quickly eliminate poor candidates and absorb risk; in the latter, it is to increase the probability of launch. This segmented approach is a good bet for your company if 60% to 80% of candidates would be eliminated at the early stage and 70% to 90% of the rest would go on to have successful market launches – and if, per project, the early stage would cost between one-fifth and one-fiftieth as much as the late stage. (Ranges are approximate; they reflect potential for a separate early stage in most situations.)



therefore, early investment in large-scale processes usually does not pay off.

While no company has replicated the Chorus approach precisely, there are examples of its principles at work in nonpharmaceutical industries. At one global chemicals company, for example, NPD suffered from both types of decision-making errors (ignoring evidence that challenged assumptions and abandoning candidates too early). To fix the problem, the company implemented carefully staged decision making, rigorous progress reviews, and strict timelines for NPD projects. But skillful project champions would invariably marshal whatever numbers and materials were needed to win support at reviews for their projects. As a result, NPD's failure rate didn't improve after the new processes were put in place. Then, recognizing the need for different

mind-sets in early- and late-stage development, management altered its recruitment strategy, working with HR to identify truth-seeking personalities for the early stage and success-seeking types for the late stage. That simple change improved NPD productivity.

In another case, a global semiconductor manufacturer - realizing that its reward systems created a disincentive for killing dicey projects early - redefined its systems to promote fast, evidence-based failure (in other words, to encourage truth seeking). This company, too, experienced improvements in NPD productivity - although, as we've noted, speed to failure is only one ingredient of successful NPD.

To gauge whether a model like Chorus's would make sense in your organization, determine whether your NPD process can be rationally segmented into early-stage development, in which you absorb risk by culling poor prospects, and late-stage development, in which you maximize the probability of launch. As a rule of thumb, in a good risk-based segmentation, 20% to 40% of all assets (such as drug candidates) or projects make it to the late stage, and 70% to 90% of those end up having successful market launches. A good segmentation also yields a per-asset cost ratio of between 1:5 and 1:50. That is, moving an asset or project through the early stage costs one-fifth to one-fiftieth as much as moving it through the late stage. (See the exhibit "When a Separate Early Stage Makes Sense.")

Consider the segmentation of drug development: If the early stage comprises Phase I and early Phase II clinical trials, and the late stage is made up of late Phase II and Phase III trials (post-POC studies), then about 20% of all candidates entering early-stage development will move on to the late stage, and about 70% of those will have successful market launches. Typically, the late-stage cost per candidate is about 10 times the early-stage cost. Thus the relationship between risk absorption and cost places pharmaceutical NPD within the bounds of good segmentation. Other industries where NPD would meet the criteria for good segmentation include biotechnology and medical devices. In industries that have a higher probability of technical success at the outset - such as cell phones, software, and consumer products in general - early POC and segmentation may do little more than extend cycle times.

Companies that could benefit from an early-stage NPD unit like Chorus need to be aware that the approach is not just a form of process reengineering. They will have to create a new, separate organization that focuses on truth seeking. A small team must be selected to plan, implement, and manage that organization. The team builds the infrastructure and recruits both internal staff and consultants, who, as discussed, may bring essential expertise and objectivity to the project. Being able to ask the right questions and design the critical experiments to rule in or rule out a product's key attributes are essential skills for people in Chorus. Teams within the unit are small and fluid, composed of individuals motivated by intellectual curiosity. Each team member works on several products simultaneously, and of course, no one will follow any of the products into later stages - a rule created to promote objective truth seeking.

As the early-stage organization develops its unique capability, it will work in parallel with the established NPD operation. It offers additional capacity but does not replace existing NPD functions. The goal for any early-stage organization and, indeed, for R&D overall should be to head off costly downstream attrition of unpromising projects. Chorus offers a promising model for reducing risk and improving R&D productivity.

Reprint R0803G
To order, see page 135.

