Production Capacity Investment With Trial Result Updates

Philip M. Kaminsky[1] Ming K. Yuen[1]

May 3, 2011

We consider the capital investment problem faced by a firm that must invest in capacity while gathering information about whether that capacity will actually be needed. Specifically, we are motivated by pharmaceutical firms, which face long horizons for clinical trials, "all-or-nothing" demand for new products depending on the outcome those clinical trials, long lead times for the construction of manufacturing capacity, and extremely high costs for construction of that capacity.

Development of a new drug typically requires a long period devoted to research, development, clinical trials, and government approvals. Indeed, it takes an average of eight years to complete the required clinical trials to prove that a new drug is safe and effective before the drug is approved by government regulators (such as the FDA in the US or the EMA in Europe). To ensure that a potentially life-saving drug is available as soon as possible, and to take full advantage of the profitable period of commercial operation while the new drug is still under patent protection, the firm would typically like to be able to produce and sell the drug immediately after it is approved. Under US patent law, for example, the firm has sole right to produce and sell the new drug under its patent for 20 years, but this patent life starts at the beginning of the initial clinical trial, so significantly less patent life is typically remaining by the time the drug is approved. Since construction and licensing of a new facility often takes four or five years, it may make sense for the firm to invest in capacity before the results of trials are known. One one hand, there is obvious risk associated with committing hundreds of millions of dollars before there is enough confidence that the drug will ultimately be approved. On the other hand, the potential losses associated with being unable to meet demand immediately can be enormous. Thus, the firm must determine when there is enough confidence to commit to building production capacity. Often, in more realistic situations, the firm has additional options – building can be paused and restarted or aborted, and there may be alternative ways to provide the capacity, with different lead times and costs.

Since the successful completion of clinical trials is the key to the success of the investment, it would be a natural step for the firm to look at the ongoing stream of information updates from the clinical trial. It is in fact considered good practice and an ethical requirement for the firm to closely monitor the progress of the clinical trial, so that the trial can be cut short if the drug performs significantly better than expected, or the trial can be shut down if the drug has adverse affects on patients. However, this information is typically not used by firms in the industry to inform capacity investment decisions. In this paper, we develop a model in which the firm actively manages investments using this information flow in order to minimize investment and penalty costs.

Model: To keep our analysis tractable, we build our model in a slightly simplified setting. We assume the firm is entering the final phase of clinical trials with enough time remaining to build necessary capacity if the trial is successful. We assume that each experiment or test case within the trial will either be a success or a failure, so that in the specific context of a pharmaceutical clinical trial, the drug is only compared to a placebo with a known success rate to determine its effectiveness. Also, the results of individual tests are revealed one at a time, on a regular basis. We assume that capacity construction projects can be

[1]Department of Industrial Engineering and Operations Research, University of California - Berkeley.

started and stopped as necessary, and feature deterministic per period building cost, setup cost when the firm decides to start or restart a project, and a deterministic penalty cost at the end of the horizon to account for lost sales and patent life.

In our model, we assume the regulatory agency requires a standard conservative "frequentist" approach for estimating the success rate of the clinical trial, in contrast to Bayesian methods that might introduce potential biases (and indeed, this is almost always the case in practice). Although the frequentist approach is required by regulatory agencies, the firm can use a Bayesian approach to estimate the likelihood of passing the trial, and use to make capacity building decisions. Indeed, the key to our analysis lies in the fact that the stochastic process we use for modeling the information update captures the Bayesian statistics which allow the firm to assess the likelihood of passing the clinical trial. This Bayesian approach allows the firm to utilize the information stream from monitoring the clinical trial, and to incorporate expert opinions and/or observations made in earlier phases of clinical trials. With ongoing observations of the clinical trial, the firm refines its estimate of the success rate, which allows us to capture the trade-off between the firm's opportunity to act early to leave investment options open, or to wait until it has better information.

Specifically, we model the capacity investment problem as a discrete time dynamic program with a finite horizon. In contrast to more standard settings, however, we assume that the transition probability of the underlying stochastic process is not time homogeneous, significantly complicating analysis. We assume that the result of a single case is revealed in each period, and we model the result of clinical case as a Bernoulli random variable with a common (but unknown) success rate, so that each case is either a success or a failure. In order to ensure satisfactory and statistically significant performance of the drug, the trial will be considered successful if the number of successful cases reaches some exogenously determined threshold over the horizon. If the clinical trial has not been sufficiently successful, the product will not be produced, and the production facility is not needed.

If the clinical trial is successful, the firm must produce the drug in order to meet demand, and to produce the drug, the firm must build a facility with sufficient capacity. In the most general version of our model, the firm can select from several possible types of facilities, each of which requires a given (not necessarily consecutive) number of time periods to complete construction, with associated per period construction costs. At the start of each period, a clinical case is completed and the firm's manager observes the result, as well as the cumulative number of successful cases up to that time, and the remaining time to construct each of the possible capacity options. The manager then decides which one type of capacity, if any, to invest in during the period. Once an investment decision is made, the corresponding investment cost is incurred. At the end of the horizon, if the product is approved, the firm pays a penalty if production capacity is not yet ready (intended to model the delay in receiving revenue from the product, loss of patent life, loss of revenue, loss of goodwill, etc.). We assume that this penalty is a function of the time needed to finish building the required capacity. The manager's objective is to minimize the total discounted expected cost of building the necessary capacity as well as any penalty cost.

Results and Analysis: We begin by analyzing the primary stochastic process in our model, and find that *the conditional probability of a successful clinical trial at any time is monotonically increasing with the number of successful results observed up to that time.* This result is the key to our subsequent analysis, in which we characterize optimal investment policies under a variety of cost structures, terminal costs, and alternative facility options.

In the most basic setting, with only one type of production capacity available, linear production costs, no additional costs associated with starting or restarting construction, and a terminal penalty cost linear in the remaining construction time, we show *the existence of a threshold level for the number of successful results observed up to that time. If the number of successful results is at or above the threshold level, it is optimal to build the facility in a given period, and if the number of successful results is below that level, it is optimal to not build, regardless of whether construction has already started.* Moreover, we show that *the decision threshold is monotonically increasing in time*, and *if it was optimal for the firm build in the previous period, it is also optimal for the firm to invest in the current period if the latest result is positive.*

We also consider a fixed cost associated with initially starting or restarting the project after idling it for one or more periods. We show *the existence of a pair of critical levels that characterize the firm's optimal building decisions in any time period. If the firm built the facility in the previous period, and the number of positive results until the current period is above the lower threshold level, it is optimal to continue building, and if the number of positive results is below the lower threshold, it is optimal to stop building. If the firm did not build in the previous period, the higher threshold level comes into play. Specifically, if the number of positive results is below the higher threshold level, the firm will start (or restart) construction in this period. Otherwise, the firm will not build in the current period.*

It is natural to relax the assumption that the terminal penalty cost is linear in remaining construction time – in particular, we would expect a long delay to be significantly more costly than a short one. Although this significantly complicates the analysis, we are able to show that *if the penalty function is increasing and convex in the number of period of construction remaining at the end of the horizon, the firms optimal strategy is still characterized by a threshold policy, and that the threshold level is decreasing in construction progress.* However, it is not true when the penalty function is not convex. In particular, when the penalty function is concave in remaining construction, we construct examples in which the action region is disconnected.

Finally, we consider a firm facing two alternative production facility options, one of which has a shorter construction lead time but higher construction cost than the other. Depending on problem parameters, the firm may find the more expensive alternative more attractive, since the firm can wait for better information before committing to build. *We identify conditions under which the firm will only select one of the two possible projects, rather than switching between the two projects as information improves.*

In a computational study, we further explore the relationship between the firm's decision between alternate projects as a function of information and problem data. We identify scenarios when the firm will interrupt construction of one project with the intention of initiating construction of the other project. In additional experiments, we explore the robustness of the firm's decisions to the reliability of initial expert opinion or results, and we analyze the sensitivity total cost to changes in problem parameters including construction and penalty costs. We also find that a larger initial sample can significantly reduce the expected total expected costs regardless of whether the trial is ultimately a success, and that if initial information is likely to be unreliable, which is often the case, a short lead time but expensive capacity option is valuable, even if it has a much higher construction cost than alternative options.

In subsequent research, we intend to introduce models that that more closely capture the nuances of clinical trials, with an ultimate goal of developing practical approaches for mitigating capital investment risk for pharmaceutical firms.