

Evidence-Based Decision Support For The Biopharmaceutical Industry

Leonard Wesley
MedLex Inc.
5941 Optical Court
San Jose, CA 95138, USA
lpwesley@medlex.com

and

Kim Graham
Cambridge Antibody Technology
3260 Hillview Avenue
Palo Alto, CA 94304, USA
kim.graham@cambridgeantibody.com

ABSTRACT

We introduce and describe an evidential-based approach to drug development decision making. The approach treats decision-related data as evidence for or against pursuing competing development alternatives. A benefit of the proposed evidential approach is that decision makers are not forced to provide data, as in traditional decision analytic approaches, that are not available or cannot be obtained from reliable sources. As a consequence, the described approach is able to give better insight into the impact of assumptions, approximations, and uncertainty on multi-million dollar decisions. A summary of results from using an evidential approach to help identify false positive and false negative results is presented.

Keywords: Decision Support, Evidential Reasoning, Portfolio Management, Pharmaceutical, Risk Management, PharmAssist™.

1. Introduction

The biopharmaceutical industry spends 15% of sales annually to discover and develop new drugs. While the total dollar investment in research and development continues to grow, the sales return for capitalized R&D has fallen to 0.85 (2003) from 2.5 (1980)¹ It has been reported that the cost and time to develop a successful drug now exceeds \$850 million and takes ten to fifteen years, while the failure rate for drug development is currently more than 90%.^{2,3} This trend in growing inefficiency cannot continue without dire consequences to the industry and to human health care.

Significant cost savings and reduced business and health risks can be achieved if weak drug candidates are identified and dropped from the development pipeline earlier and faster than is currently the practice. In a published report, The Boston Consulting Group states⁴,

“...If a company were able to decide in just one out of ten ... cases against pursuing a target in the first place, it would save as much as \$100 million per drug ... As for INDs that fail clinical trials, if the company was also able to decide in just one out of ten cases to abandon development earlier, it could save an additional \$100 million per drug.”

While there are many reasons for the high cost and failure rate of drug development, it is clear that even a marginal improvement in deciding which drug candidates to develop (or not develop) can result in a huge ROI and reduced risk.

A recent market survey revealed that biopharma perceives the greatest challenge (i.e., “pain”) to making effective drug development decisions is identifying when and the extent to which assumptions, approximations, and guesses can impact multi-million dollar decisions.⁵ The same survey revealed several additional challenges to reducing the risks associated with making decisions that include

- the lack of a structured and systematic process by which the variability in the uncertainty associated with making particular decisions can be quantified and tracked
- identifying the decision parameters that warrant closer scrutiny and deliberation
- accelerating and facilitating cross-functional and inter-department decision making

Biopharma's current approach to addressing these issues varies between company, department, function, and corporate philosophy. Some, such as the previous Smith Kline Beecham, now part of GSK, have advocated, since the late 1990s, a more formal process-oriented and traditional decision analytic approach to decision making.⁶ Some business and finance departments within biopharma companies base drug development decisions, in part, on calculations from well known asset valuation formulas such as NPV, ENPV, and so forth.⁷ Yet others use methods and approaches ranging from using commercial decision software such as CrystalBall and DPL through in-house developed software and spreadsheet calculations.

Companies continue to try to systematically evaluate the scientific, economic, political and corporate factors necessary to make optimal business decisions. To assist with this difficult process, an increasing number of companies are starting to employ decision-support methodologies that are based on traditional probabilistic techniques. These techniques are used to help determine next steps on everything from advancing lead compounds, in- and out-licensing, selecting clinical trials to pursue through managing the corporate portfolio.

Some users of traditional decision analytic and probabilistic technologies are perhaps aware of the requirements and assumptions that the calculus requires, but they struggle to satisfy them. Unfortunately, many are unaware of the requirements and assumptions one must adopt before credible results can be obtained.

In practice, when decision makers are forced to provide data that are lacking or are of unknown reliability they will often adopt approximations or unwarranted assumptions or even make guesses to fill in the missing data. This can lead to poor analytic results that reflect a convergence of good and speculative data, imperfect assumptions, guesses, and approximations. These analytic results in turn influence decision makers to pursue, or not pursue the development of drug candidates. Poor analytic results can cause what might be called decision "blind spots" with respect to discerning promising from weak drug candidates. When multi-million dollar decisions hinge on having accurate probabilistic estimates, then it behooves one to understand these requirements and assumptions, and to become aware of sound and proven state-of-the-art technologies that can address important issues related to these and other aspects of decision-making.

Here, we describe a state-of-the-art "evidence-based" approach to making decisions that can be employed at any and all stages of the drug develop process. The term "evidence" follows from the fact that the data and information used to make decisions cannot always

"pinpoint" the optimal decision. Rather, often a more appropriate view is that such data are best treated as evidence that tends to support or refute, to varying degrees, probabilistic arguments for choosing between competing alternatives. The approach described here exploits this notion in a mathematically sound and proven manner, and within a computational environment that allows one to readily and intuitively

- quantify the degree to which assumptions and approximations might impact decisions
- identify which decision parameters deserve the greatest or least attention
- facilitate a structured cross-functional and inter-departmental approach to decision making.

2. The Probabilistic Challenge

We begin describing the requirements for improved decision making by summarizing some of the requirements for using traditional decision analytic and probabilistic approaches.

It is well know that traditional probabilistic approaches to decision making require providing complete statistical data before a probabilistic assessment can be produced. In other words, the probabilities of every possible value of each factor (or parameter) that is deemed relevant to discerning which drug development alternative is most appropriate to pursue must be obtained, derived, computed, or approximated. For example, suppose future sales and market position at the time a drug candidate is approved are two of many important and complex factors for assessing the probability that a drug candidate merits continued development in a company's pipeline. Further suppose that the possible values for future sales are low, medium, and high, and that the possible values for market




Figure 1: Example probability distribution for the sales factor.

position are first, second, and third. Before a probabilistic estimate can be produced, the probability that future sales will be low, the probability that sales will be medium, and probability that sales will be high must be specified. This probability specification is commonly called a *probability distribution (pd)* that is illustrated for this

example in Figure 1. A *pd* must be specified for market position as well as is the case for all remaining factors being considered.


In practice, the complexity of the real world precludes us from always knowing which *pd* should be used for each factor. The probability of future sales of a drug typically depends on many things. These include, but are not limited to, future population, disease, and income demographics as well as the political climate and the cost of drugs to the consumer. The influence of these factors as well as other factors must somehow be captured in the *pd* that is used to specify the required probabilities. Such *pds* are extremely difficult, if not impossible, to calculate. What do people do if they cannot specify all of the required *pds*? Some will guess at a distribution to use, or use *pds* that were used in the past. At times a uniform distribution will be used to convey total ignorance about the probability estimates for each possible value of parameters such as sales, see (Figure 2a). Others might use a normal distribution if it is believed it will reflect the probability of a future market position (see Figure 2b).

Indeed there are an infinite number of distributions one can choose from, but only one is the correct distribution. Any other distribution will be inaccurate to varying degrees. Consequently, guessing at which *pd* to use will in all likelihood result in selecting a non-optimal *pd*. One might argue that it is sufficient to use a *pd* that is “close enough” to the true *pd*. How can we assess “close enough” in the absence of the true *pd* or if previous “similar enough” situations do not exist that would allow extracting reliable statistics? Of course, if we had the true *pd*, there would be no need to make guesses or approximations! In practice, we will likely never know the true distribution of all complex and critical factors.


Later we will describe how a particular calculus that is employed by a commercial product called PharmAssist™ can quantify the degree to which guesses or assumptions about probabilistic estimates might impact multi-million dollar decisions. PharmAssist™ can also help decision makers identify which parameters to focus on to lower risks.

One alternative to making educated guesses at selecting a *pd* is to use Monte Carlo simulation techniques to derive the desired probability distribution. Unfortunately, even Monte Carlo methods require specifying a *pd*, or in the case of factors that range over a continuous numeric scale, a *probability density function (pdf)* for each modeling parameter. Which *pdf* or *pd* should be used? Once again, if the existing distribution, (normal, uniform, triangular, and so forth) adequately specifies the probability of a factor’s possible values, then an accurate result can be obtained. Once again, in practice, most will pick “some” distribution and trust that the drug development decision that will be made based on the

results from guesses or approximations made in either Monte Carlo simulations or other probabilistic methods will not result in an undesirable influence on decisions.



(a)



(b)

Figure 2: (a) Example uniform probability distribution for the sales factor. (b) Example normal distribution for market position factor.

Another well known method that can be used to help deal with uncertainty about a *pd* or *pdf* is the use of second order probabilities. Second order probabilities are probabilities about probabilities, sometimes called meta-probabilities. One specifies a second order probability, or a distribution for each *pd* that one is trying to specify for each decision parameter. Analytic calculations would result in not only a probabilistic assessment of the decision alternatives, but a “second order” probabilistic estimate that characterizes the likelihood that the assessment of each decision alternative is correct. Using second order probabilities, however, simply defers the inevitable need to determine a meta-probability number or meta-*pd* about the original *pd* for parameters.

We believe a more desirable approach to making drug development decisions is, from the start of the decision making process, to not force users to provide statistical data that are not available or from credible sources. Rather, we propose using analytic technologies that can provide accurate results about what is and is not truly known. In the following section, we introduce

PharmAssist™ and the notion of an evidential-based approach to decision making.

3. An Evidential Approach

The origins of the topic of evidential reasoning and an evidential-based approach to decision making can be found in the upper and lower probabilistic work of Arthur Dempster⁸ and the belief function (BF) work of Glenn Shafer⁹. This work formed the basis of evidential reasoning (ER)¹⁰.

A belief function primer:

Within the BF calculus, possible answers to questions like “What should be done with the current drug candidate?” are represented as a set of interrelated propositional statements, called a *frame of discernment* (FOD) or simply a *frame*. Elements within a frame represent a mutually exclusive and exhaustive set of propositions that can be viewed as possible answers to some question of interest. For example, a discrete “DRUG CANDIDATE” FOD $\Theta = \{DROP, SHELVE, OUT-LICENSE, ADVANCE\}$ might be used to represent possible answers to the above question. FODs can also be represented on a continuous scale, say from 0 to 100. Propositions in a continuous representation might be one or more sub-intervals of the interval $[0,100]$. For example, the interval $[34,56]$ in a continuous DRUG CANDIDATE FOD might correspond to the proposition “*The potential of the current drug candidate to address the disease of interest is between 34 and 56.*” The choice between a continuous or discrete FOD depends on the form of the answer that is desired by the user. For simplicity, the discrete FOD above will be used to introduce the belief function calculus without loss of generality.

Discerning the true proposition in a FOD involves acquiring and combining opinions about the truthfulness or falsity of subsets of Θ from distinct sources. Sources of information can range from statistical data obtained from previous results, through subjective estimates of things like the potential therapeutic index, and competitor’s progress on the same mechanism.

In practice, a single body of information will not always directly suggest the truth or falsity of propositions in a FOD. Rather, input information (called an opinion) is often expressed in terms of propositions in a frame that is distinct from and indirectly related to possibilities in the DRUG CANDIDATE frame Θ . Before a consensus can be formed about the true proposition in the DRUG CANDIDATE frame, opinions that are specified in indirectly related FODs must be first translated to the DRUG CANDIDATE frame. The relationships between propositions in different frames are represented as a compatibility relation (CR) that is defined to be

$$CR : 2^{\Theta} \mapsto 2^{\Theta} \quad \text{Eq. 1}$$

where every mapping from $p \in 2^{\Theta}$ to $q \in 2^{\Theta}$ means that if p is true then it is possible that q can be simultaneously true, and vice versa. The notation 2^{Θ} refers to the power set of Θ and represents all possible combinations of elements of Θ , including the universal (everything) and the null set (nothing, \emptyset).

To illustrate this, consider the DRUG CANDIDATE and REGULATORY HURDLES FODs and compatibility relation that is shown in Figure 3. The presence of a link between propositions in distinct FODs means that the propositions can be simultaneously true or simultaneously false. In this example, if the Regulatory hurdles are “High”, then it is possible (not necessarily probable) that the potential of the current drug candidate is “Low” or “Medium”, but not “High”. Similar interpretations can be made for other links. The absence of a link means that the compatibility relation does not support evidence that the propositions can be simultaneously true or false.

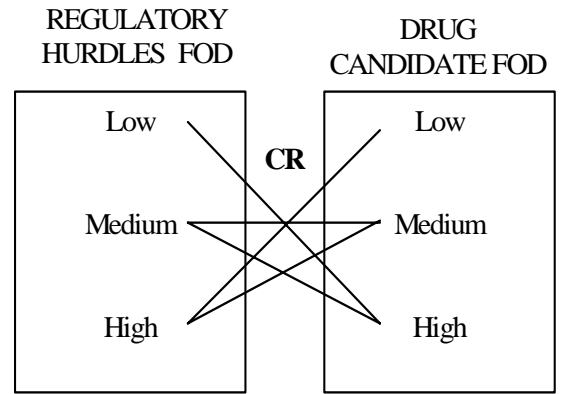


Figure 3. Example compatibility relation between the REGULATORY HURDLES and DRUG CANDIDATE FODs.

To assess the likelihood that propositions in the DRUG CANDIDATE FOD are true or false, opinions about the truth or falsity of propositions in related FODs, such as REGULATORY HURDLES, TOXICITY, and so forth must be obtained and translated to the DRUG CANDIDATE FOD. Once these opinions have been translated to the DRUG CANDIDATE FOD, they are combined into a composite estimate using Dempster’s rule⁷.

To accomplish this, each source of an input opinion has a unit of mass m that can be attributed to subsets of Θ based on the source’s perceptions and beliefs about the truth and falsity of possibilities. The mass m is distributed as follows:

$$m(\emptyset) = 0, \quad \sum_{A \subseteq \Theta} m(A) = 1, \quad \text{and} \quad Bel(A) = \sum_{\substack{p \in A \\ A \subseteq \Theta}} m(p) \quad \text{Eq. 2}$$

where Bel represents the total mass (or Belief) in support of the proposition based on the available evidence. One might be 80% certain that the regulatory hurdles are high. This opinion can be represented by attributing 80% of the unit mass to the “high” proposition in the REGULATORY HURDLES FOD. If a user remains uncertain about other possibilities, this can be conveyed by attributing the remaining 20% of the unit mass to the disjunction ({Low \vee Medium \vee High}) of all the propositions in the REGULATORY HURDLES FOD. Note that this represents a genuine indication of lack of knowledge rather than the assumed distribution among the different options as dictated by traditional probabilistic approaches. Similarly, one might be 60% confident that the level of TOXICITY is also high. The remaining 40% uncertainty can be attributed to the disjunction of all propositions in that frame. The number of “input” FODs and compatibility relations depends on the type of information and relationships among them and the DRUG CANDIDATE FOD that one desires to model.

Once opinions over the desired related FOD's have been conveyed, they can be “translated” via compatibility relations to the DRUG CANDIDATE FOD.

For example, suppose we are 60% certain that the regulatory hurdles are High. Further suppose that we are 80% certain that the TOXICITY of the drug candidate of interest is also High as shown in Figure 4. After translating both opinions through the compatibility relations CR1 and CR2 to the DRUG CANDIDATE FOD, a consensus can be formed using Dempster’s rule. Dempster’s rule is used to reach a consensus between two opinions as follows:

For all $B_1, B_2, B_3 \subseteq \Theta$,

$$m_3(B_3) = (1-k)^{-1} \sum_{B_1 \cap B_2 = B_3} m_1(B_1) m_2(B_2),$$

where

$$k = \sum_{B_1 \cap B_2 = \emptyset} m_1(B_1) m_2(B_2) \leq 1$$

Eq. 3

and k represents the total conflict between m_1 and m_2 . Dempster’s rule can be repeatedly applied to the result of a previous combination (e.g., m_3) and the next mass distribution, if any, to arrive at a new consensus. The rule is also multiplicative in nature, which means the order of combination is immaterial and can be carried out in parallel to the degree supported by computational hardware.

Opinions that are conveyed, translated, and combined in this manner induce what is commonly called an evidential interval within the belief function calculus. An evidential interval (EI) is associated with every proposition (e.g., “High” potential) $A \subseteq \Theta$. The lower (Spt) and upper (Pls) bounds of the interval represent the

minimum and maximum degree of support that can be attributed to a proposition given the current opinions, and is defined to be $[Spt, Pls] \subseteq [0, 1]$ where

$$Spt(A) = \sum_{\substack{A \subseteq \Theta \\ p \in A}} m(p), \text{ and } Pls(A) = 1 - \sum_{\substack{A \subseteq \Theta \\ p \in \neg A}} m(p) \quad \text{Eq. 4}$$

The interpretations of evidential intervals are as follows:

- [0,0]: evidence suggests proposition is completely false
- [1,1]: evidence suggests proposition is completely true
- [n, 1]: evidence suggests proposition is partially true, for $0 < n < 1$
- [0, n]: evidence suggests proposition is partially false, for $0 < n < 1$
- [n,m]: evidence suggests proposition is partially true and partially false, for $0 < n < m \leq 1$
- [0,1]: totally ignorant about the truth or falsity of the proposition

A benefit of using the BF calculus to assess the drug candidate’s potential is that the EI associated with each possibility in the DRUG CANDIDATE FOD allows one to distinguish the total amount of information (or evidence) that suggests the corresponding proposition is true from the degree to which we are ignorant about the drug candidate’s potential. This measure is captured in the width of the EI; that is, the $Pls - Spt$ of a proposition.

With this distinction, the use of sensitivity analysis techniques can be used to help determine which factors, such as REGULATORY HURDLES or TOXICITY, have the greatest or least impact on the drug candidate’s assessed potential. This provides an opportunity to identify and address the areas impacting a drug development decision, thereby improving the quality of the decision or providing guidance on the areas that require the most immediate attention before final decisions are made.

In Figure 4, the mass in the Risk FOD corresponding to the translation of both the REGULATORY HURDLES and TOXICITY FODs is specified by the translation operation. With these two results now expressed in the same FOD, Dempster’s rule can be used to combine them and form a consensus opinion. In this case, the two translated mass distributions only “agree” with respect to the “Low” potential proposition. Dempster’s rule accentuates the level of agreement in the final mass distribution as shown in Figure 4. The evidential interval of [0.6,1.0] indicates that the decision maker is justified to at least degree 60%, and no more than 1.0, that the potential of the drug candidate is “Low”. A more colloquial way of interpreting this interval is that the pessimist would believe the potential is “Low” to a degree of 0.6 because this is the amount of evidence that supports the proposition, while the optimist would believe the “Low” proposition completely true because there is no evidence that contradicts it. The difference

between these two is a measure of the uncertainty that the potential is Low.

Comparison to Traditional Probabilistic Approaches To Decision Analytics:

There are several distinguishing aspects between traditional probabilistic and statistical inference schemes and the belief function calculus. One aspect is with respect to the underlying assumptions and requirement that the sum of the support attributed to disjoint elements of a space of propositions must be unity. This is reflected in the mass distribution function m discussed above. Quite often when statistical data are lacking many decision makers will use a uniform distribution to convey their ignorance about the likelihood of parameter values. The problem is that it might not be the case that the true probability of propositions is uniform, yet the traditional probabilistic approach forces us to "pick a probability value."

Another requirement is the additive law or probability where the $\text{Prob}(A) = 1 - \text{Prob}(\sim A)$. At times, it might not be warranted to assume, in the absence of evidence, that the $\text{Prob}(\sim A)$ is 0.8 just because there is evidence to suggest the $\text{Prob}(A) = 0.2$. In other words, traditional

probabilistic and statistical methods require adopting assumptions and requirements that are sometimes undesirable or are awkward to deal with when trying to make complex decisions. This requirement of "the traditional math," we posit, can at times lead biopharma to make inappropriate drug development decisions.

An attractive aspect of the BF and ER calculus is that if and when complete statistical data happens to be available, Dempster's rule gives the same analytic result as Bayes' rule produces. In that sense, BF and ER can be viewed as a generalization of traditional probabilistic methods.

4. PharmAssist™

PharmAssist™ is a commercially available decision analytic and portfolio management software technology that embraces the ER calculus.¹¹ Unique to the PharmAssist™ platform is the ability to say, "I don't know" for a given input factor. PharmAssist™ can carry uncertainty forward through the analysis process in a mathematically sound and justifiable manner. After an analysis is completed, an explanation of the results can be requested and displayed so that the user can obtain guidance about the factors that have the greatest and least influence on the presented results. The user then has the option to obtain the additional information needed to refine and reduce the uncertainty reflected in analysis results. Once the uncertainties are resolved, the confidence of the decision increases. If one decides not to obtain better information, at least decision can be made with better knowledge about which factors have the greater potential to negatively impact decisions.

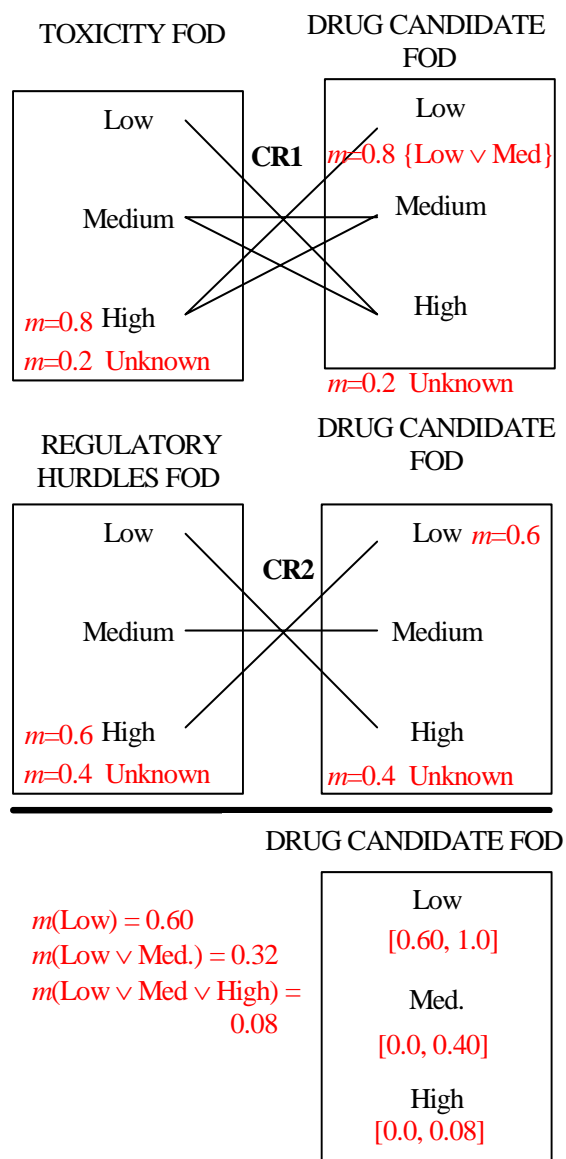


Figure 4. Example DRUG CANDIDATE FOD with evidential intervals that reflect a consensus of uncertain opinions about the potential due to regulatory hurdles and TOXICITY.

PharmAssist™ also employs sensitivity analysis to assess the degree to which assumptions, and approximations, and guesses impact decisions. PharmAssist™ provides the results of such sensitivity analyses in a manner that can be displayed as illustrated in Figure 5. In a simple and intuitive manner, the assumptions that matter are the ones with the greatest disparity between the analysis with assumptions and the analysis when the assumptions were removed by "saying I don't know" as described above.

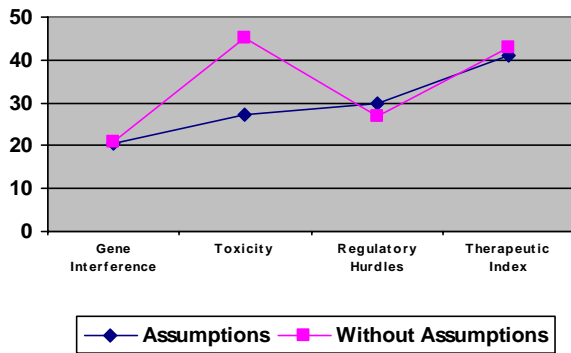


Figure 5: Example graph of PharmAssist™ results that indicate which assumptions are having an impact on analytic results. The parameters exhibiting the greatest disparity between the case when assumptions are and are not made are the parameters of interest.

5. Results

A retrospective analysis of drug development decisions has been conducted on several well know successful and failed drugs, one of which was a cardiovascular agent that was initially approved and launched and then withdrawn after a few months. This analysis involved obtaining and inputting FDA data that were available before the decision under analysis was made. In other words, no data beyond the decision point were input to avoid PharmAssist™ taking into account data “after the fact.” The results demonstrated that PharmAssist™ was able to identify drug toxicity and the duration of clinical trials as two parameters that warranted significant deliberation before moving the above-mentioned cardiovascular drug into clinical trials. This case is an example of PharmAssist™ catching a false-positive. Additional results revealed that PharmAssist™ is capable of identifying additional false positives and a fewer number of false negatives.

6. Conclusions And Future Work

The biopharmaceutical industry is well aware of the need to address its growing inefficiency on several fronts. Better technologies for screening and assessing new compounds, more and better assay tests, and so forth. Among these is a recognition of a need for better technologies for managing the assumptions, approximations, uncertainty, and risks that are inherent with making complex drug development decisions.

We have described an evidential-based approach to managing decision-related information that removes the need for decision makers to provide statistical information that is not available or from credible sources. As a consequence, the proposed evidential approach to making drug development decisions provides decision makers greater insight into how their assumptions might impact multi-million dollar decisions than is currently possible using traditional methods.

Some logical next steps for PharmAssist™ include extending the use of ER techniques and algorithms to multivariate analysis, and to provide decision recommendations given corporate objectives, constraints, and requirements.

REFERENCES

- ¹ SG Cowen & Company, “Biotechnology Industry Overview,” March 2005.
- ² Ashburn, Ted T., and Karl B. Thor, “Drug Repositioning: Identifying and Developing New Uses For Existing Drugs,” *Nature Reviews: Drug Discovery*, Vol. 3, pp. 673-683, August 2004.
- ³ Singh, Mahesh, “Successful Portfolio and Resource Management Practices for Pharmaceutical & Biopharmaceutical Companies: Industry Benchmark Report,” Published by Pittiglio Rabin Todd & McGrath, 10275 West Higgings Road., Suite 410, Rosemont, IL 60018, Phone: 847-813-4542, (December 2003)
- ⁴ Tollman, Peter, et al., Boston Consulting Group, “A Revolution in R&D: How Genomics and Genetics are Transforming the Biopharmaceutical Industry, November 2001
- ⁵ “Decision Support Market Study,” Sponsored by MedLex Ins., Conducted by US Market Access Center, 111 North Market Street (6th Floor), San Jose, CA 95113-1101, December 2005.
- ⁶ Sharpe, Paul, and Tom Keelin, “How SmithKlineBeecham Makes Better Resource-Allocation Decisions,” *Harvard Business Review*, March-April 1998, Reprint # 98210.
- ⁷ Kellogg, D., & J.M. Charnes, “Real-Options Valuation for a Biotechnology Company,” *Financial Analysts Journal*, pp. 76-84, May-June 2000
- ⁸ Dempster, A.P., “Upper and lower probabilities induced by a multivalued mapping,” *Annals of Mathematical Statistics*, Vol 38, No. 2., pp 325-339, (April 1967).
- ⁹ Shafer, Glenn, *A Mathematical Theory of Evidence*, Princeton University Press, 1976.
- ¹⁰ John D. Lowrance, Thomas Strat, Leonard P. Wesley, Thomas D. Garvey, Enrique Ruspini, David Wilkins, “The Theory, Implementation, and Practice of Evidential Reasoning,” Final Project Report for DARPA, (SRI Project 5701, Contract No. N00039-88-C-0248, Task B), (June 1991).
- ¹¹ Developed and marketed by MedLex Inc. www.medlex.com