I&EC research
Industrial & Engineering Chemistry Research

pubs.acs.org/IECR

Article

Gambling on Innovation

Darrell Velegol*

Cite This: *Ind. Eng. Chem. Res.* 2021, 60, 7689–7699

Read Online

ACCESS | Metrics & More | Article Recommendations | Supporting Information

ABSTRACT: The purpose of this article is to provide a method for choosing an innovation portfolio, which is based on a simple and fundamental result from information theory. Given a set of innovation projects, estimated payoffs and probabilities of success for each, and an established capital fund, the question is "How much should one invest in each project in order to maximize the growth and reduce the risk of going bust?" That is, how should innovation leaders gamble on their innovation projects? The question is part of the process design for innovation processes. The concept in this article is especially important for early-stage innovation projects, where the probability of success is often <50%. As a heuristic, for $p = 50\%$, one needs a payoff of $b = 4$ to achieve a 25% median compound annual growth rate. This article provides an algorithm to guide investment decisions in innovation, plus some heuristics that can be used for guidance. It also emphasizes the critical importance of estimating both the probabilities of success and the payoffs for projects.

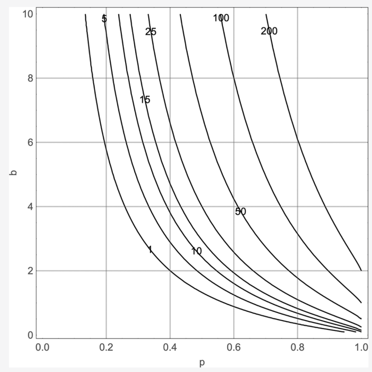

1. INTRODUCTION: THE QUESTIONS IN THIS ARTICLE

The race is not always to the swift, nor the battle to the strong, but that's the way to bet.—Damon Runyon.

The purpose of this article is to provide a practical method for choosing an innovation portfolio, which is based on a fundamental result from information theory. Given a set of innovation projects, estimated payoffs and probabilities of success for each, and an established capital fund for innovation (i.e., an initial innovation bankroll), the question is "How much should one invest in each project in order to maximize the growth and reduce the risk of going bust?" That is, how should innovation leaders gamble on their innovation projects? Equation 1 and Figures 5 and 7 give the key results.

The allocation of innovation resources is an essential decision for chemical research to reach fruition and is part of the process design for innovation processes. A common approach is to base the allocation on the estimated net present value (NPV) and/or internal rate of return (IRR) or similar commonly used measures for the projects. One estimates the investment required over time and the revenue expected over time, discounts these back to the present-day value (e.g., with the cost of equity, CoE, as the discount rate), and takes the difference. If NPV > 0 with IRR > CoE, the project is a "yes" and otherwise a "no". However, this approach suffers from three important challenges. (1) Not every innovation project results in a success, perhaps failing at the R&D level, the commercial level, the regulatory level, or others. Not all companies estimate a "probability of success" for their payoffs, which can be disastrous. It is like playing blackjack without knowing the odds. As this article shows, it is essential to have an estimate of the probability of success. For early-stage innovation, oftentimes $p < 50\%$. (2) The NPV–IRR style approach still does not indicate what fraction of our initial innovation bankroll we should invest in each project. In the extreme, should we simply pick out one, two, or three highest NPV projects and put all our investment there? Or as many as we can afford until we have spent out? Or some other strategy? (3) As we will see in this article, the NPV or IRR types of "arithmetic average" approaches inevitably lead to "going bust" over time. Your company might still limp along, like an engineer who loses every weekend at the casino but remains solvent due to a steady income, but the innovation gambles that you are making are losing money, or they are at the least inferior to what they could be.

The core concept behind this article is a method familiar to the investing and gambling communities, the Kelly criterion (KC). In 1956, shortly after Claude Shannon had published his famous article on information theory,[1] Kelly sought to use the ideas of information theory to improve performance in games of chance. He wanted to find the maximum growth rate in total wealth for a gambler with a private but potentially noisy wire of information. As Kelly stated in his article, "The maximum exponential rate of growth of the gambler's capital is equal to the rate of transmission of [Shannon] information over the channel."

In fact, we might see the role of a company's innovation team (including R&D, as well as commercial leadership, marketing, manufacturing, legal, finance, regulatory, safety, and other functions) as providing information that reduces the risk (i.e.,

Received: February 3, 2021
Revised: April 29, 2021
Accepted: April 29, 2021
Published: May 17, 2021

probability of a non-success) of an innovation idea failing somewhere in the process. There are systematic methods for reducing this risk and increasing the speed of innovation.[2,3] In the academic world, one usually publishes an article only with a high probability—perhaps >99% or even 99.9%——that the work is correct. However, attaining a 1.0 or 0.1% probability of failure (risk) is very expensive in terms of time and money, and in a competitive marketplace, a manager needs to know how to allocate investments in part to avoid being scooped. It is well-known in the investing world that asset allocation is among the most important factors of success, that is, choosing when to pull money from one class of investment and put the money elsewhere including even cash. In this article, I provide a quantitative method for doing this with innovation investments, and as we will see in Section 6, the method can be extended to a broad range of activities.

To introduce the concept, let us explore an example, which will provide some intuition on how to place our bets on innovation projects. Let us say that I enter a coin flipping game at a casino. The casino lets me use a coin from my own pocket, which I believe is unbiased, such that the probability (p) of heads is 0.5 and tails is 0.5. In this game, if heads comes up, my payoff odds are given as 1.5 (i.e., if I bet \$1 and win, I increase to have \$1 + \$1.5 = \$2.5; so here I define a payoff ratio b = 1.5), and if tails comes up, I lose my bet (i.e., if I bet \$1 and lose, I now have \$1 less; I define a loss ratio a = 1). The casino offers me the opportunity to make 1000 flips, and if I start, I must finish all 1000 flips or forfeit any winnings. If I start with \$100, how much should I bet each round? The bet is clearly biased in my favor, as the casino knows; however, there must be some reason why they offered the bet. Wanting to maximize my wealth at the end of the 1000 bets, I don't want to squander my advantage. So should I bet it all? Before I do so, thankfully I recognize that putting "all my eggs in one basket" is probably not wise either. Note that this game is not an ergodic process. Placing 1000 simultaneous bets at one time period is much different from placing 1000 bets consecutively and independently.

We can simulate the result. What happens if I bet 34% each round? Starting with \$100, I'll bet \$34 on the first flip. If I win, I now have \$100 + \$34 × 1.5 = \$151. On the second round, I would bet 34% of my new amount, or \$51.34. If I continue this pattern for all 1000 flips, then simulations show that my median outcome would be about \$18 left after 1000 flips, and I would in fact go bust ("ruin", where I reach <\$1) about 41% of the time. Aha! That's why the casino made me this generous offer! And if I were to bet 50% each time, I would go bust >99.9% of the time, and my median final wealth would be \$0. That is, by betting 34 or 50%, I'm putting "too many eggs in one basket", again even though I have the clear advantage in the betting. My arithmetic average outcome will be even higher but only because some rare runs make it so; in fact, my median outcome is awful in these cases. This is the reason why financial investors use a portfolio of (hopefully independent) investments rather than putting the entire investment into equities for instance.

If, by contrast, I were to bet at half of the previous rate, or 17% of my capital each round, my median take home amount after 1000 bets would be about \$72 B, and I would go bust (<\$1) with a probability of <0.0065% (i.e., once every 154,000 trips to the casino, or almost never). How could I possibly know to bet "so little" to win so much? This is the problem that Kelly solved in 1956.[4] The Kelly criterion in eq 1 below, derived in Supporting Information Section A, in fact gives f_{KC} = 0.1667. Plugging in p = 0.50 (probability of winning) and $q = 1 - p = 0.5$ (probability of losing), with b = 1.5 and a = 1 gives f = 0.1667. By maximizing the growth rate of the total wealth, he established what is now known as the KC for the fraction of your wealth (f_{KC}) to gamble in a binary (win−lose) bet

$$f_{KC} = \frac{p}{a} - \frac{q}{b} \tag{1}$$

In eq 1, for losses, a = loss divided by the bet amount (i.e., the investment, dimensionless, often with a = 1), and for wins, b = gain divided by the bet (again, still dimensionless, although we can have $b \gg 1$).

Thus, the KC gives the bet fraction that maximizes the growth rate of the capital *and* gives a very small chance of going bust. Any other bet size, either higher or lower, has a sub-optimal growth rate and gives a smaller final capital.[5] It has been hypothesized that in some of the loss aversion experiments done by Kahneman, Tversky, Thaler, and others, we are not so much seeing "loss aversion" as we are seeing "ruin aversion" given many successive chances.[6]

Below are several questions that a gambler—or an investor, an innovation director, a CTO, an assistant professor, a private equity professional, or a federal agency manager—might want to know in order to maximize the total wealth while maintaining a low probability of gambler's ruin (i.e., close to zero wealth, which in practice might be defined as 1 or 10% or another % of the original). This article addresses questions 1−4 below; questions to be answered in future articles are given in Section 7. In my experience, for most companies, these questions remain unmeasured, or ad hoc rules of thumb are used, usually without quantitative justification. However, the KC concept provides a way to assess more rigorously and profitably. Of course, we must realize that low probability events do occur: In your industry, there might be entrepreneurs betting all they have on ideas that would disrupt your company, and on the slim chance that they are right, they might be David to your Goliath. Sometimes, Hail Mary passes are successful in football. Sometimes, players win at roulette. However, these low-probability gambles are not the way to bet consistently, as the Runyon quote says. Here are the questions I answer in this article:

1. Allocation of bets. What fraction of my initial innovation bankroll (W_0) should I bet on each of my potential innovation projects? Which bets should I avoid entirely?
2. Quantiles of the compound annual growth rate (CAGR). If I have a set of bets, each with a binary success probability (p) of payoff (b) and a probability of loss ($q = 1 - p$) of amount (a), what is the anticipated median CAGR? Is it greater than the cost of equity (CoE)?
3. ruin rate. How often will I "go bust"? We could choose any fraction to define "ruin", but here, we will define "ruin" as losing 99% of your initial investment. Your company might still feed a bad innovation process, keeping it afloat, but the portfolio selection might make it a loser.
4. Algorithm and heuristics. Is there a simple and practical algorithm that I can use to allocate my portfolio of bets? Are there guiding heuristics that I can use in the absence of more detailed knowledge?

Despite its advantages, there are two well-recognized shortcomings of betting according to the KC.[7] (1) Finding good bets. This article provides a method for evaluating known opportunities, but it does not provide a route for identifying or generating new opportunities.[3] (2) A relatively high early allocation. While it is true that the KC maximizes your long-term

https://doi.org/10.1021/acs.iecr.1c00511
Ind. Eng. Chem. Res. 2021, 60, 7689−7699

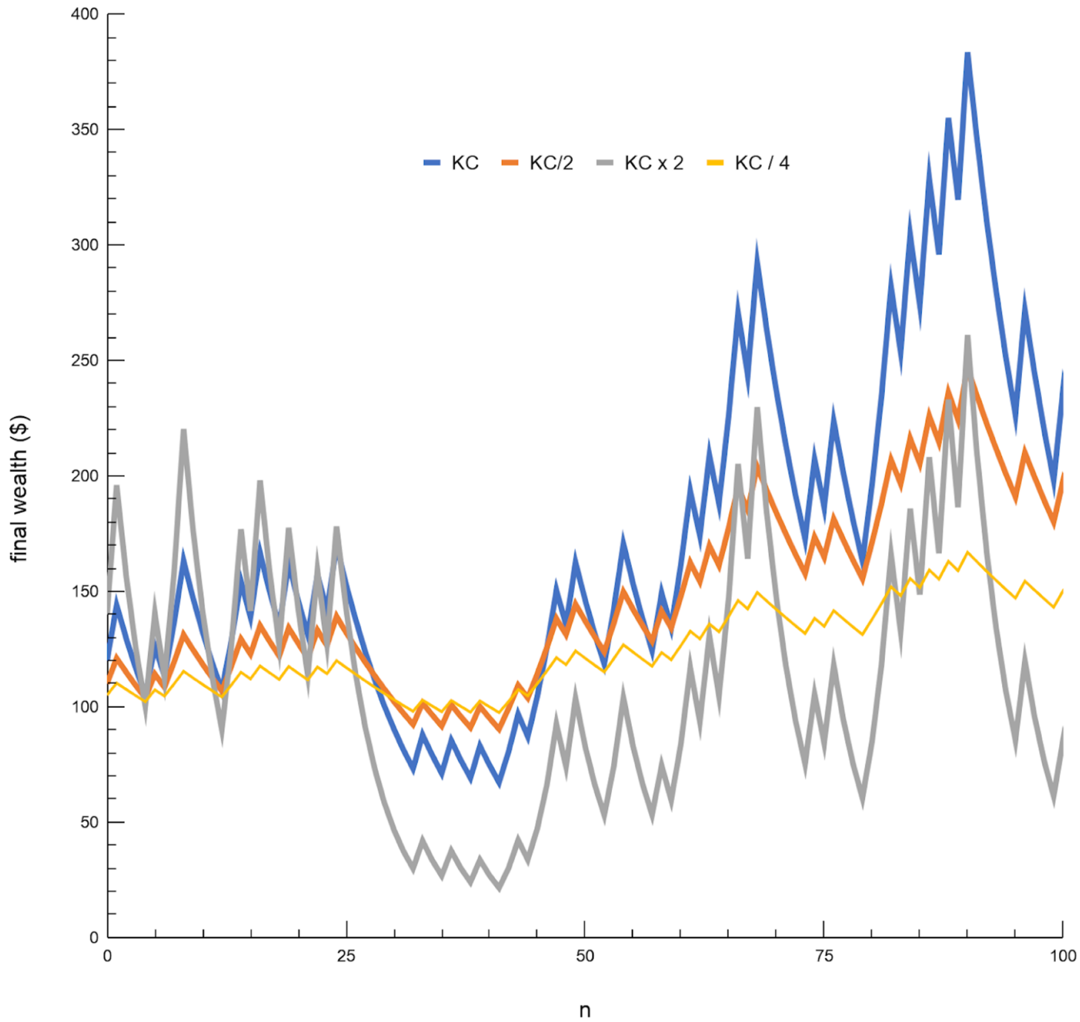

Figure 1. Total wealth as a function of time up to n = 100 independent betting points with b = 2 and p = 0.4 for an initial wealth of $100. The KC fraction f = 0.10. Note the high volatility in the outcome. I show the Kelly bet of eq 1 (KC), the fractional half Kelly bet (KC/2), the quarter Kelly bet (KC/4), and the double Kelly bet (KC × 2). The double Kelly bet gives the lowest value in this simulation, although it does not quite "go bust". The KC bet gives the highest bankroll at n = 100, followed by the half KC and then quarter KC bets. This is one particular run, but there tends to be high variability for one run such that it is not entirely uncommon for the half KC bet to give a final W higher than that of the full KC. However, the half KC bet always has less volatility than the full KC bet, and the quarter KC has even less volatility. If we were to extend play every alternative history that could have occurred in this simulation of n = 100, we would find that the median final W = $264 for KC, $209 for KC/2, $154 for KC/4, and $107 for 2KC.

growth rate, the initial allocation is still volatile, and so, many investors or gamblers avoid going bust early by using a "fractional Kelly bet", often half. I point out that there are critics of the KC for investing, perhaps most notably, the late Nobel laureate Paul Samuelson.[8] His primary critique was that maximizing the growth rate is equivalent to maximizing a logarithmic utility function, but that there are other utility functions. Ziemba wrote a helpful article,[9] not disputing Samuelson as much as showing how his arguments sharpen the theory and its effectiveness, for instance, by showing that "there are no guarantees". Sometimes, a high probability of winning can still leave us broke, and sometimes, a long-shot wins the race.

2. ALLOCATION OF BETS

In this article, we have a binary outcome: Your "innovation bet" either wins or loses. You have a payback ratio (b) with a probability (p), and a loss ratio (a) with a probability ($q = 1 - p$). Here, I note two important points about p. (1) The value of p depends not only on R&D, but the whole innovation chain, including marketing, manufacturing, legal, regulatory, safety, and other functions. (2) The value of p will likely change as expenditures are made. The role of "learning" will be assessed in a future article.

Given this, for question 1, the KC gives a quick answer: From eq 1, if $f > 0$, you should bet the fraction indicated. If you have multiple bets, bet each according to their fraction, for any $f > 0$. There are in fact situations where the sum over all projects $\sum f_i$ might be higher than 1 (e.g., $\sum f_i = 1.18$, as in Section 6), and if you are able to leverage and borrow money, then the KC suggests to do so. If the capital available is truly fixed, you need to use the numerical method outlined in Section 5. If $f < 0$, you should avoid the bet or even "sell short", although how to do this might be non-obvious for investing in R&D. This might mean just nixing the project, but perhaps, this means selling the technology to another organization while retaining the option to buy it back if sufficient developments have been made, of course with some price for the option.

https://doi.org/10.1021/acs.iecr.1c00511
Ind. Eng. Chem. Res. 2021, 60, 7689−7699

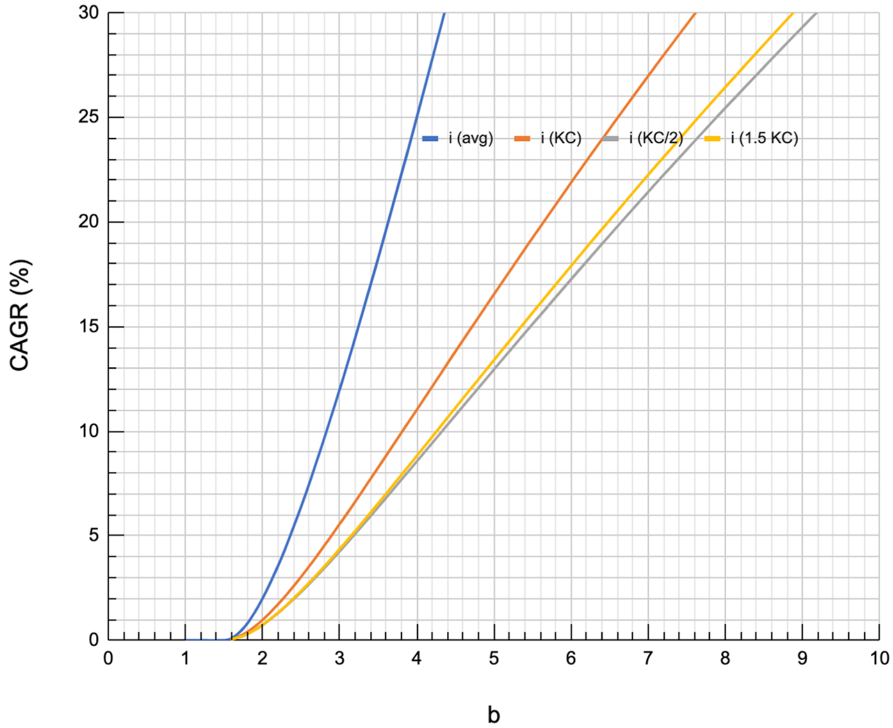

Figure 2. CAGR (here, labeled as i for various values of b, with p = 40% probability of success and bet sizes of multiples of the KC, 0.5, 1.0, and 1.5 × KC. For b < 1, the KC recommends not betting since after many trials, you face nearly inevitable losses for p = 0.40. The blue curve is the arithmetic-average CAGR using the KC bet from eq 1. This average is skewed by some large but rare outcomes. The orange curve shows the median geometric-average CAGR using the KC bet. We see that the fractional "half KC bet" (KC/2) gives about 75–80% of the median CAGR; however, the half KC bet has a ruin rate of < 0.1%, while the full KC bet has a ruin rate of up to about 2%. We see that for this p = 0.40, we need b > 4.5 to get a CAGR of 15% with the KC bet and b > 6.25 to get a CAGR of 25%.

The key insight that Kelly had is that gamblers are usually *not* seeking to maximize the arithmetic mean of the final wealth, averaged over all possible outcomes. The arithmetic mean is skewed by some rare but very large outcomes. The *arithmetic-average* final wealth (W, unit \$) over n bets, given by the ratio ($\bar{\mathfrak{R}}$) over the initial wealth (W_0, units \$), is

$$\bar{\mathfrak{R}} = \frac{W}{W_0} = (1 + fpb - fqa)^n \tag{2}$$

This assumes that all n bets are of the same size and probability. The result is readily generalized using a product if each bet is different.

It turns out that if $pb - qa > 0$ (i.e., NPV > 0), then not only does eq 2 indicate to bet your entire bankroll (f) on the gamble to increase the average growth but it also says that you should leverage as much as you possibly can (i.e., start with the maximum bankroll W_0). However, as I will show, oftentimes, the arithmetic average goes up even as the median wealth goes to zero,[10] and your most likely outcome becomes ruin. That is, you have "put too many eggs in one basket". As we will see in the section on "ruin", betting a value (f) of roughly double f_{KC} leads to almost inevitable long-term ruin.

Alternatively, Kelly recognized the importance of maximizing a *geometric average* for final wealth, which can be written as

$$\mathfrak{R} = \frac{W}{W_0} = (1 + fb)^{pn}(1 - fa)^{qn} \tag{3}$$

Equation 3 says that we grow for each time we win ($W = pn$) and diminish for each time we lose ($L = qn$) for the n games. This Bernoulli-style equation applies for an ongoing activity of consecutive and independent events rather than saying that we both win some amount and lose some amount on each bet as in an arithmetic average. For both eqs 2 and 3, the CAGR (given by i) is given by

$$i = \exp\left(\frac{\ln \mathfrak{R}}{n}\right) - 1 \tag{4}$$

which gives the arithmetic-average CAGR

$$i = f(pb - qa) \tag{5}$$

and the geometric-average CAGR

$$i = (1 + fb)^p(1 - fa)^q - 1 \tag{6}$$

As we see in eqs 4–6, a higher value of n does not increase the CAGR (i), but since that growth rate occurs over more attempts, eq 3 shows that the final wealth is larger for a higher n. This fact would recommend faster innovation processes to achieve a higher n in a given time.[3] As shown in Supporting Information Section C, for a revenue (R), a profit margin (M), and an innovation investment (E), we have

$$b = \frac{RM}{E} - 1, \qquad a = 1 \tag{7}$$

Equations 6 and 7 together give the result we seek. If we can estimate the revenue, the profit margin, and the expenditure for an innovation investment—for instance, by using a modified Delphi method[11,12]—then we can calculate the fractional return on the investment (b). As discussed in Section 5, if expenditures (E) and revenues (R) are done over time (i.e., the usual case), then we must discount them all back to $t = 0$ (NPV). Then, we

https://doi.org/10.1021/acs.iecr.1c00511
Ind. Eng. Chem. Res. 2021, 60, 7689–7699

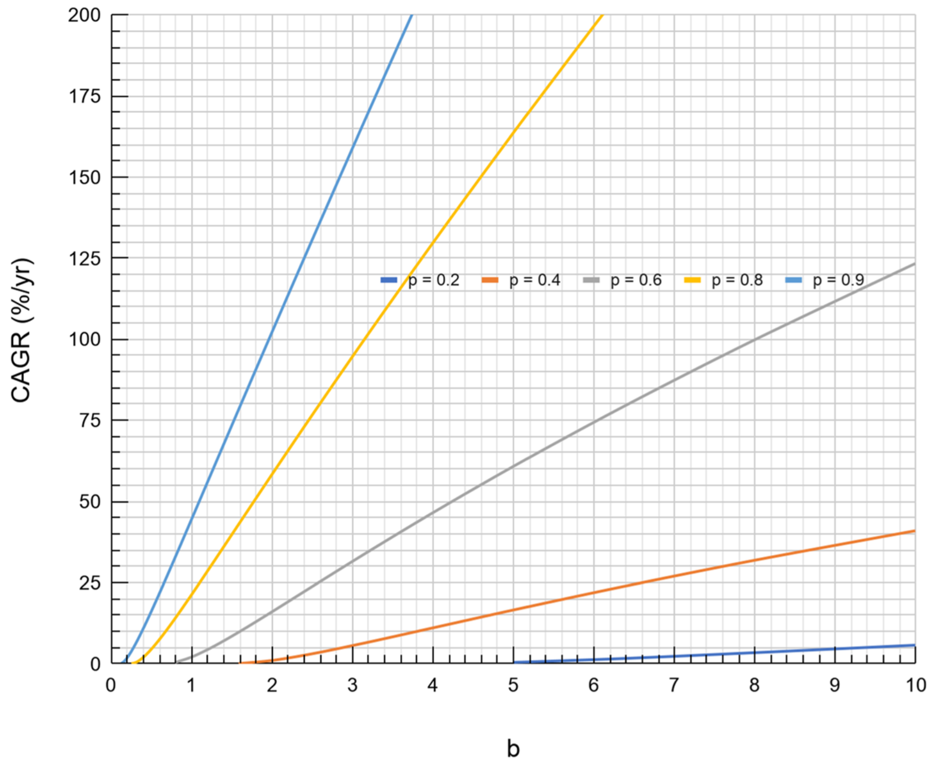

Figure 3. Median CAGR (%/yr) for various values of p and b for the Kelly bet $f = f_{KC}$. When the probability of success is low (e.g., $p = 0.2$), the KC indicates that only a large value of $b > 5$ makes betting profitable, although perhaps not even worthwhile. As seen in the graph, the results become roughly linear for large b; for example, when $p = 0.2$ and $b = 100$, the CAGR is 53.8%, compared with 5.69% for $b = 10$. We thus see two key points: (1) to "bet high payback" (i.e., high p and b) and since the CAGR is compounded with the number of bets (n), also (2) to "bet often", which means to have innovation processes that make attempts faster.

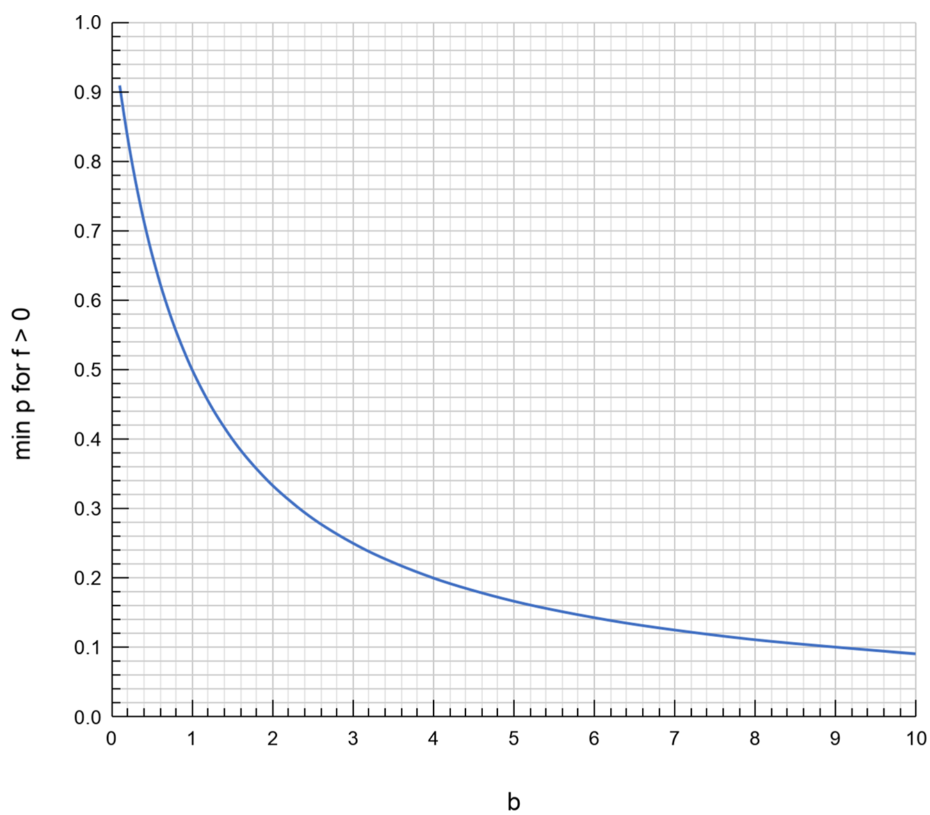

Figure 4. Minimum probability of success to have a viable gamble from eq 8. The curve is based on $f = f_{KC} = 0$ and uses $a = 1$. If you are below the curve, repeated gambles on those bets will lead to loss in the long run since $f_{KC} \leq 0$. For $a = 1$, eq 8 gives $p = 1/(1 + b)$, which is the same as this plot.

https://doi.org/10.1021/acs.iecr.1c00511
Ind. Eng. Chem. Res. 2021, 60, 7689−7699

use the KC exactly as in eq 1 to get the Kelly bet. To assess the CAGR, we then use eq 6.

3. QUANTILES OF CAGR

We assume that our betting process is Markovian, meaning simply that the next result depends only on the current position, not the history of all that has come before. This is not entirely true for innovation investments since even for unsuccessful outcomes, there are some learnings that might add value to future investments. This learning possibility will be considered in a subsequent article. The details for the calculations in this section are given in Supporting Information Section B. In short, we approximate the Bernoulli approach in eq 3 with a normal distribution.

Figure 1 shows a simulation using random numbers for p = 0.40, b = 5, and a = 1. The starting investment is \$100. The KC bet rate goes above \$383 at one point. The fractional KC bet (often done using $f = 0.50f_{KC}$) is less volatile than the full KC bet, and the quarter KC bet is less volatile still.

In Figure 2, I plot the median CAGR (i) for various values of b with p = 0.40, a specific value. For these curves, the ruin rate when we use the KC value for f is about 2%. The ruin rate here is calculated as the fraction of trials for which the final amount is <1% of the original amount, which is not \$0 but very bad. The KC gives a return that increases with b at p = 0.4 and avoids ruin. The return is the highest for the Kelly bet, but interestingly, when $f = f_{KC}/2$, the CAGR is still about 75−80% of that when $f = f_{KC}$. That is, for a lot less risk or volatility, the return is still quite high. In Figure 3, I plot the median CAGR for various values of p and b for the Kelly bet $f = f_{KC}$. I recognize the importance of seeking higher p and b based on the figure, and Figure 3 implicitly tells us to seek a higher value of n in the same time (i.e., faster innovation) since the CAGR is compounded by the number of bets (n).

From Figure 2, we see that the KC provides a superior betting strategy compared with either higher or lower bet fractions (f). We might want to know what is the minimum probability that we should have to place any finite bet $f_{KC} > 0$ for a given a and b. Solving eq 1 for $f = 0$ gives p as a function of the payoffs a and b.

$$p = \frac{\frac{1}{b}}{\frac{1}{a} + \frac{1}{b}} \tag{8}$$

Figure 4 shows a plot of the minimum p needed for a given value of b for a = 1.

The KC shows that if b = 1, then to achieve f_{KC} = 0, you need a probability of success of at least about p = 0.50. Thus, if b = 1 and p = 0.4, you should not enter that bet; it is a long-term loser. This sets a fairly high bar for innovation returns! If your b = 0.5, then you must have p > 0.7 in order to bet fruitfully over time. Of course, individual bets might pay off even with worse probability values—after all, people do win at roulette or at the horse track—but this gambler's approach is a long-term losing strategy for investment. We can in fact assess the combinations of b and p that will give any CAGR we choose, and these values are given in Figure 5, a key figure in this article.

4. RUIN RATE

This section looks at two important questions concerning "ruin". (1) What value of $f > f_{KC}$ will cause me to go bust and (2) how often will I be ruined (i.e., in this case, be left with <1% of the original innovation bankroll) if I bet a certain fraction (f) on a

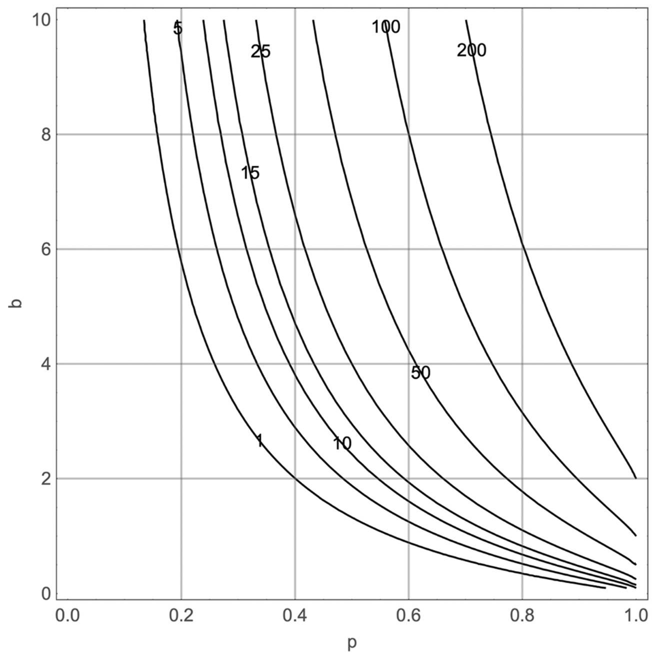

Figure 5. Contours of constant CAGR (from 1 to 200%) for combinations of p and b, betting the Kelly fraction f_{KC}, which is the best you can achieve. To achieve a certain CAGR, you can trade off b and p. For instance, you can achieve a 15% median CAGR by having approximately either {b = 8, p = 0.3} or {b = 2, p = 0.6}. Thus, one can exploit low p bets if the b is sufficiently high. The CAGR shown in this figure should be higher than your cost of capital (or equity) to make the investment worthwhile.

certain bet? The answer to the first of these questions is given in Figure 6 for particular cases of p and b: Roughly, when $f > 2f_{KC}$, going bust is almost inevitable in the long run (since the median CAGR < 0). A corollary from Figure 6 is that if you have uncertainty in the values of p and b, you will likely predict a sub-optimal value for the fraction (f) to bet, and your CAGR will suffer. Since the plot is concave, errors will always be harmful to profits. Figure 6 enables us to calculate the "value of perfect information" to assess whether it is worthwhile to spend money improving our estimates of p and b in order to increase profits.

For the second question, Figure 7 gives the ruin result for several values of p. As the plot shows, for very low values of b, the ruin rate is 0% but only because the KC indicates not to place a bet. Generally, the lower the probability of success (p), the higher the ruin rate. Once again, it leads us to consider that a key purpose of an innovation process is to reduce the uncertainty and so increase the value of p, which reduces the amount of unknown information entropy.

Here, as in almost all of our analysis, we see the importance of evaluating b and p for our projects. In my experience, this is too seldom done with accuracy in most organizations. A key point is that higher values of p lead to lower rates of ruin. As said earlier, a core task of innovation is to increase the values of p and lower the risk of the project. It is like having a private wire for a gambler, in Kelly's words.

5. PRACTICAL METHOD FOR ALLOCATION AND REALIGNMENT

As in all investments, the allocation of investments (to *achieve* investment diversity) and rebalancing over time (to *maintain* investment diversity) are among the top priorities of the

https://doi.org/10.1021/acs.iecr.1c00511
Ind. Eng. Chem. Res. 2021, 60, 7689−7699

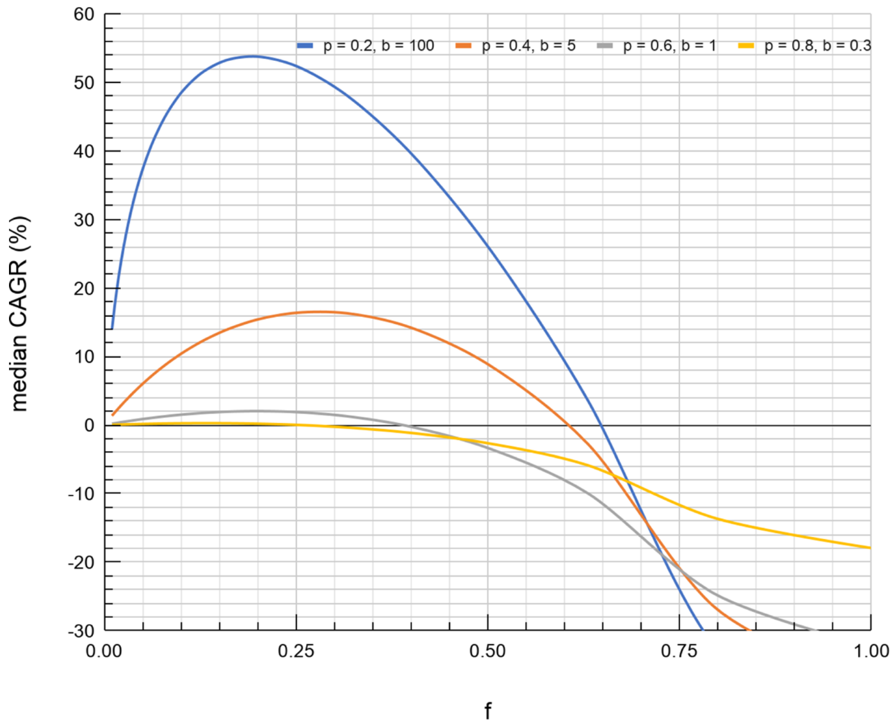

Figure 6. Median CAGR has a maximum with f. The plot is shown for the four cases listed later in Table 1, with given values of p and b. As one moves away from $f = f_{KC}$, CAGR falls off roughly quadratically, so that small errors in f—including those due to imperfect estimates of p, b, and a—are costly. For the high $b = 100$ curve, the median CAGR crosses 0 at $f/f_{KC} = 3.39$. For the other cases ($b = 5$, 1, and 0.3), this happens at a ratio of 2.18, 1.75, and 1.88, respectively.

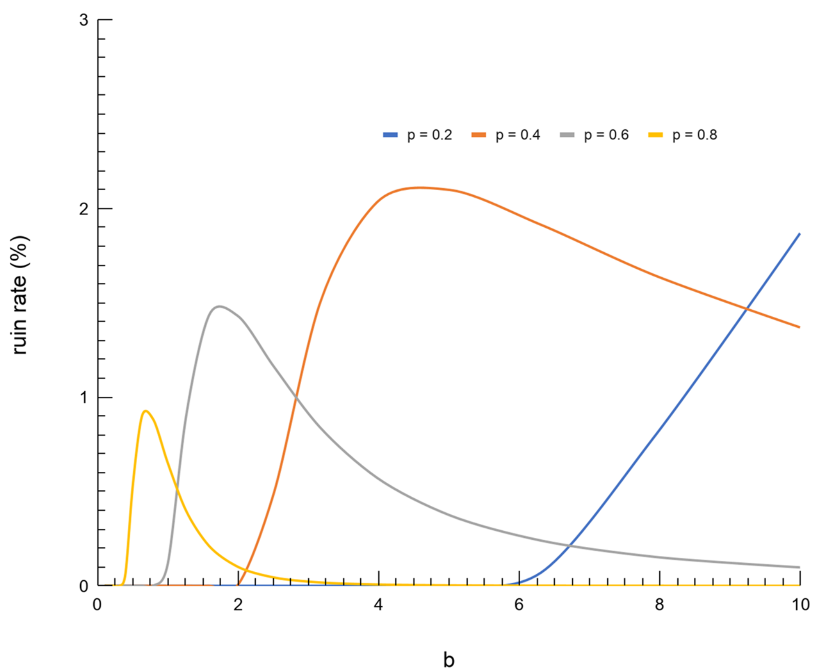

Figure 7. Ruin rate for various values of p and b using the KC bet of eq 1. For a small b, each p has a ruin rate of 0% but only due to the fact that the KC indicates not to place the bet. The lower the probability of success (p), the higher the ruin rate.

investor. Below is an algorithm for choosing your innovation investment allocation. Since in this article we do not yet consider time from multiple perspectives, we do not yet require that we choose a time frame (T) over which you want to maximize your rate of growth. This is an important topic for a subsequent article.

1. Assess the potential payoff (b), loss (a), and probability of success (p) (and therefore failure probability $q = 1 - p$) of all projects available. List uncertainties if possible, or at least recognize that there is uncertainty in these numbers.
2. Assess the total amount of capital you have to invest in innovation or the amount of expenditures (W_0) you have available to invest in innovation (units \$) for the year.

https://doi.org/10.1021/acs.iecr.1c00511
Ind. Eng. Chem. Res. 2021, 60, 7689−7699

Table 1. Fraction of Allocation (f) and Investments for the Example Given in the Text[a]

project #	b	p	a	$q = 1 - p$	$f = f_{KC}$ (eq 1)	invest ($MM)	multiplier	ruin %	mean ROI %
1	100	0.2	1	0.8	0.192	1.92	1.5382	0.0107	369
2	5	0.4	1	0.6	0.280	2.80	1.1654	0.1414	39.2
3	1	0.6	1	0.4	0.200	2.00	1.0203	1.4895	4.00
4	0.3	0.8	1	0.2	0.133	1.33	1.0028	0.0204	0.53
					0.805	8.05	1.8341		683

[a]We see that the total investment is less than the anticipated \$10 MM. We hold \$1.95 MM in cash, rather than investing it, to maximize our growth rate. In this case, the final median wealth would be \$18.34 MM. If we sought to maximize the arithmetic mean, we would invest \$100 into project 1, and the final mean wealth would be \$202 MM; however, we would almost surely (>99%) go bust. Using the fs calculated from the KC, the arithmetic mean is shown in the right hand column. Multiplying all the $(1 + i)$ gives 6.83, meaning the final wealth would be \$68.3 MM "on average" (arithmetic) over all possible outcomes.

3. Identify your weighted-average cost of capital (WACC) or CoE. Remember, this is not determined by you but by the external market.
4. Calculate the investment fractions $f = f_{KC}$ for each project based on the KC in eq 1. If you prefer less volatility in your investments, but with 75−80% of the benefit, invest at $f = 0.5 f_{KC}$ (i.e., "fractional Kelly"). Remember to calculate b using the revenue discounted back to $t = 0$ with your own WACC.
5. If the $\sum f_i = 1$, then you have your allocation. More likely, the fs will not sum to unity. If $\sum f_i > 1$, you have two choices. If you are able to leverage your position with $\sum f_i > 1$, attain the extra funds and do this. If $\sum f_i$ must be 1 or less, then you need to solve numerically as given below this numbered list. As a heuristic, having $\sum f_i > 1$ means you should bet more toward your innovations with a higher multiplier, usually meaning a higher anticipated value of b.
6. If $\sum f_i < 1$, you can either return the extra money to Central as un-needed surplus, or you can negotiate to hold this money as cash for which you will strike a new investment when a good opportunity arises. This is similar to holding cash as part of your personal investment portfolio. You might have the understanding that you would return to Central any unspent funds at the end of the year as long as next year's allocation would not be decreased. Note that this requires great trust and discipline not to squander the funds on projects for which $f < 0$, just to avoid "use it or lose it". Use it or lose it is a harmful heuristic for companies, federal funding agencies, and governments at all levels. It is better to hold the funds in cash than to gamble on innovations with $f < 0$. Yes, sometimes these can win, just as people win at roulette. However, that is not an effective way to bet.

When the $\sum f_i \leq 1$ is a hard constraint (i.e., no leveraging possible), the following set of equations must be solved numerically (e.g., either Mathematica or Excel Solver works well). Following eq 3 but extending to multiple bets gives the following

$$\max \Re = \frac{W}{W_0} = \prod_{i=0}^{N} (1 + f_i b_i)^{p_i} (1 - f_i a_i)^{q_i}$$
$$\text{s. t.} \qquad 0 \leq f_i \leq 1$$
$$\sum_{i=1}^{N} f_i \leq 1 \tag{9}$$

Solving eq 9 gives the set of f_i values. When $\sum f_i \leq 1$, the numerical solution will give the same result as the f_i calculated using eq 1, which assumes that each bet is independent. However, when $\sum f_i > 1$, the results will begin to diverge from the independent result, favoring those bets with higher median payoffs, as expected from Figure 3. If we wanted to do so, we could add a constraint to the optimization problem posed in eq 9 to provide a ceiling on the ruin rate or a floor on the IRR. Thus, we could optimize the growth while maintaining a ruin rate less than some specified amount.

6. EXAMPLES AND APPLICATIONS

Let us look at a couple examples. Say that a company has budgeted innovation (not just R&D, but the entire innovation chain) at \$10 MM for the coming year. We have four potential projects (just four to keep this example tractable) in which to invest, and we have estimates for p and b that we trust are close to correct. For this article, we use a binary result: If we do not succeed, we lose that investment, so $a = 1$. This assumption will be relaxed in a later article. We have one project that we estimate has a large payoff of $b = 100$, with $p = 20\%$ probability of success; one with a low payoff $b = 0.30$, but with $p = 80\%$ probability of success; one with $b = 5.0$ and $p = 40\%$; and one with $b = 1.0$ and $p = 60\%$. How much should we bet on each project? Table 1 gives the result, with the fs calculated from eq 1 and the CAGR multiplier calculated from eq 6.

We have four points here: (1) We might wonder, "Where does the size of the opportunity show up in this calculation? Shouldn't the result depend upon whether I'm aiming for a \$100 MM outcome or a \$100k outcome as with an NPV calculation?" It is automatically scaled into the parameter $b = \text{RM/E} - 1$ from eq 7. In this equation, we account for the profit (revenue times margin, RM) compared with the expenditure (E). (2) Based on the KC, we are not planning to invest the full \$10 M, and so, we leave \$1.95 MM unspent as cash. This is almost unheard of when managers use a "use it or lose it" mentality. However, it is more profitable to hold cash in these circumstances. Some managers might feel "stupid" holding onto cash instead of "investing it", but as this article shows, we gain a higher median return by holding cash than investing when $f < 0$. Companies like Apple and Berkshire Hathaway are well-known for holding large cash reserves at times. (3) We see that even projects with a low b can be "favorable bets" that help maximize the rate of growth (i.e., the criterion of the KC in eq 1). The alternative to spending everything or using these low b bets is to leave the money in cash and to hold as cash until more favorable opportunities arise. (4) We see that projects with a low probability (e.g., project #1) but high payoff (b) are shown to be viable bets ($f > 0$) by the KC. Even for $p = 0.01$, we find $f > 0.4$. However, the low p bets will

https://doi.org/10.1021/acs.iecr.1c00511
Ind. Eng. Chem. Res. 2021, 60, 7689−7699

Table 2. Fraction of Allocation (f) and Investments for the Example Given Above in the Text[a]

project #	b	p	a	$q = 1 - p$	f_{KC} (eq 1)	f_{KC} (eq 9)	invest ($MM)	multiplier
1	100	0.2	1	0.8	0.192	0.186	1.86	1.5380
2	5	0.4	1	0.6	0.280	0.268	2.68	1.1652
3	1	0.6	1	0.4	0.200	0.166	1.66	1.0197
4	0.3	0.8	1	0.2	0.133	0.018	0.18	1.0007
5	4	0.5	1	0.5	0.375	0.361	3.61	1.2497
					1.180	0.999	10.00	2.2853

[a]We see that the total investment is now $10 MM. The low multiplier projects (e.g., #2 and #4 from Table 1) diminish.

have much greater volatility, and so if the leadership cannot stomach this variability, we might hold a mix of low p and high p innovation investments.

Now as an additional example, let us add one more potential project to the previous mix. We now have project 5 with $b = 4$ and $p = 0.5$. How do we allocate our investments now, if we must hold the total investment to $10 MM? For this problem, the sum of $f = f_{KC}$ is > 1, and so we must solve numerically according to eq 9. The results are given in Table 2.

Of course, part of the allocation has gone to the new option, project 5. This has come primarily by moving the previous cash holding to option 5, but we also see that the investment in project 4 (i.e., the lowest multiplier of wealth, with a payoff not much different from holding cash) has dropped from 0.133 to 0.018 in favor of other high payoff options. The total wealth will be multiplied by 2.2853 instead of 1.8341 from before adding project 5. I emphasize again the following "Who knows? Could option 4 produce a higher multiplier?" Yes, just as I could win at the roulette wheel. However, it is not the way to bet. You could win money; you could lose money; but the median multiplier is 1.0007. As a heuristic, if $\sum f_i \leq 1$, then move your excess investment toward cash. On the other hand, if $\sum f_i \geq 1$, move your money toward the higher multiplier options, as given in eq 6.

There is an additional corollary to Table 2. At the beginning of the budgeting process for resources (time, effort, and money), one might be tempted to "maximize your gains" by planning to invest all available resources. However, this implicitly assumes that no "great opportunity" like project 5 above will become available as time goes by. In my experience, such opportunities seem to arise periodically, such that an emergency need comes up for some resource—perhaps chemical analysis, another function or a piece of capital equipment. Just as many investors hold back part of their money in cash so that if a great opportunity arises, they have liquid assets so that they can strike quickly, *a successful strategy may be to plan yearly goals or milestones, leaving some capacity for emergencies or high payoff opportunities.*

We can extend the concept here to other domains. How should federal agencies invest in research? There is another nonlinearity at work here: confidence to start a project because we know it will last 3 or 4 years. Introducing this into the equations starts to swell the number of parameters we would use, and so this article would have as its main recommendation that the agency seek an estimate of b (payoff) and p (probability of success) and make its innovation wagers based on the Kelly Criterion, if their goal is to bring maximum payoff for taxpayers. For agencies with many projects (e.g., NIH), big data might start to inform the choices, showing the value of horizontal connections (i.e., including of course R&D and also manufacturing, marketing, regulatory, and other functions) rather than simply vertical research silos. The data here would include the probability of various types of projects being commercialized and the amount of the payoff for society. Agencies might explicitly ask for these estimates and their methods of estimation.

If you are a funding agency making your claim to the Congress that your program is worthwhile, and you have projects that give a 10% CAGR at 50% probability, 10× CAGR at 10% probability, and 100× CAGR at 1% probability, how should you allot your money? The KC would suggest 4.5% toward the 10% CAGR, 1% toward the 10×, and 0.01% toward the 100×. What about the rest of the money? The KC suggests you spend it on better projects, or, even better, find ways to increase the probability of the higher CAGR numbers, which I believe can be done using better innovation processes.[3] Where can we use the ideas in this article in the academic world? Imagine that you are a new assistant professor, and you have a fixed number of hours each week to invest in your projects. How do you allot these hours? What fraction do you put into project A, B, C, and so forth? Or do you put "all your eggs in one basket"? The Kelly criterion suggests the answer in eq 1, if we can establish a common utility function for all the costs and benefits. Here, we will call utility "free energy", and we will let it be dimensionless (in thermodynamics, we would divide by kT). Then, the researcher could decide how to translate hours into free energy, or winning a grant into energy. If it takes 100 h to write a proposal, you have four ideas, and you believe that project A's proposal will result in a grant p = 30% of the time, then you have the start of the process. Perhaps for taking 100 h to write, you assign an investment cost of ΔG = +4. However, perhaps winning the grant provides $\Delta G/\mathrm{kT} = -30$ free energy back to you. The payoff $b = (30-4)/4 = 6.5$. The KC in eq 1 would suggest you spend 19% of your time on this task.

If you are a venture capital firm supporting new ideas, you can estimate probabilities and payoffs to indicate what fraction of your total investment wealth you should allocate into each venture. Additionally, you can determine whether you have the ability to raise the p and b values for the company's innovation and thus increase its value quickly and dramatically. A key challenge, which again I will take up in a later article, is how to estimate the probabilities and payoffs. These estimates are notoriously bad for most companies.

For college admissions or grad school admissions, you can decide the archetype of the student and the probability of a particular outcome, convert the costs and benefits to a utility function, and calculate your own f values. One might say, "That's implicitly what we do now." However, why do it implicitly? Find better metrics so that you can compose a proper utility function, and make quantitative, explicit decisions. "Yes, but this is the real world. Those ideas don't work." Tell that to Thorp and other investors and gamblers who have done very well using such quantitative manifestations of these ideas.

https://doi.org/10.1021/acs.iecr.1c00511
Ind. Eng. Chem. Res. 2021, 60, 7689−7699

Table 3. McKinsey Innovation[a]

	familiar with tech	not familiar with tech
not familiar with market	product line extensions, new markets,on-top margin = 0−10%, <IRR≥ 20−25%, success rate = 30−40%, t = 2−7 years (avg 5), $b \approx 8$, f = 0.27	new product in new markets, on-top margin = 0−60%, <IRR> = 8−12%, Success rate = 15−20%, t = 8−19 years (14), $b \approx 12$, f = 0.11
familiar with market	product line extensions, existing markets, on-top margin = 0−5%, <IRR> = 18−23%, success rate = 40−50%, t = 2−5 years (4), $b \approx 4$, f = 0.31	new product launches in existing markets, on-top margin = 0−10%, <IRR> = 13−18%, success rate = 30−40%, t = 6−15 years (11), $b \approx 7$, f = 0.26

[a]In their article, they formulated a 2-axis scheme, the y axis with "familiarity with the market", and the x axis with "familiarity with the technology". Using middle values (e.g., for the upper left, IRR = 22.5% and success rate = 35%), we can estimate b values using Figure 5. These are given in the table. Note that the average companies in this McKinsey study have $b \geq 4$ for all quadrants.

The KC could be applied to allotting supercomputer time, NSLS X-ray beam time, and other scarce resources. One might say, "But we could never know those utility functions." It is true that establishing the utility function might be a challenge, and we might even change the function from time to time, but foregoing the KC criteria seems like a dereliction of duty for the possible gains in outcome.

I look at one more point based on an article from McKinsey.[13] They established a four-quadrant classification with market familiarity (yes−no) and technology familiarity (yes−no). Based on their interviews with R&D leaders and their own Innomatics database, they stated the IRR, the success rate, time to success, and the "on-top" percentage (i.e., the additional margin that innovations gained vs the products they replaced, net of cannibalization). Their results are summarized in Table 3 from several of their tables. By using their success rate as our probability of success (p) and their IRR as the CAGR, we used Figure 5 to estimate the equivalent payoff (b) and then the KC bet f_{KC}. These are given in the table.

Accounting for all four quadrants, the sum of the f_{KC} values is 0.95, close to unity but a bit less, indicating that a company having a portfolio with this mix of projects is investing approximately correctly. If projects in these quadrants were all used, they would multiply profits by 1.78 (calculation not shown, but based on eq 6). However, the time required for the projects is long. A decrease of 10% in the time required—quite possible in my experience—would allow 10% more projects to finish, increasing the number (n) of bets and therefore increasing the profits significantly since n is in the exponent.

There is one more consideration that I want to raise, which draws us into an econochemistry perspective. The reaction of ideas into an innovation has an activation energy cost. Without sufficient investment, the ideas might never move forward at all. In molecular reactions, if two atoms approach along the reaction coordinate with insufficient energy (e.g., at a low temperature) to overcome the transition state energy, the attempt fails and the atoms separate again. That is, the initial kinetic energy was insufficient to cause a reaction. In order to achieve a reaction, we can either raise the temperature (i.e., raise the initial kinetic energy of the atoms) or add an appropriate catalyst (i.e., make the energy required for the reaction lower). If we raise the temperature for an exothermic reaction, we decrease the equilibrium constant and make the reaction harder to proceed. A catalyst, on the other hand, enables a new pathway. In considering the KC, I might extend this concept to avoid overspending on an innovation project to "get over the transition hump", if the KC indicates not to bet or to bet small. Instead, we might find simpler questions and experiments with a lower activation energy that catalyze the probability of success so that the KC fraction f_{KC} becomes larger. Finding these more focused questions and hypotheses and doing the smaller experiments is the hard work of research, to open new pathways for ideas to react to form innovations.

7. CONCLUSIONS

Key ideas from this article are as follows: (1) there is a necessity of estimating the payoff (b), loss (a), and probability of success (p). (2) The NPV, IRR, and other "arithmetic mean measures" can lead to ruin, whereas the *geometric-mean-based* KC leads to the maximum growth rate and a low probability of "going bust". (3) Once we have a, b, and p, the method is scalable and readily useable by eq 1, with the CAGR given in Figure 5 and ruin rate given in Figure 7. The Kelly method applies to ongoing activities that have consecutive, independent events. (4) To increase the CAGR, one needs to have higher p and b (upside) and lower q and a (downside). And increasing n increases profits for CAGR > 0. Since having $p < 1$ tells you that some things are going to fail, you might as well do it quickly and wisely (i.e., "Intelligent Fast Failure").[15]

From this article, we have focused on a few critical questions and have left some others to subsequent articles. Here is an additional list of questions to be considered in future articles:

5. Roles of time and learning. How does time impact innovation investing, beyond discounting for the NPV, in a dynamic environment? Four additional places where time enters include (a) getting scooped if you are too slow, (b) learning, as less plausible ideas and hypotheses are removed and new discoveries are made, thus making a, b, and p change with time, (c) innovating faster, so that a higher n is achieved in the same time, and (d) avoiding the specified use of all resources at the beginning of the quarter or year so that if an emergency or high payoff opportunity arises, you have liquid resources available to exploit the opportunity.
6. Estimating p and b. The evaluation of these critical parameters is not trivial, and errors or uncertainties can be translated into dollars from the ideas in this article. What if we know $p(b)$, that is, a probability distribution over a range of payoffs (b)? Or even better, what if we know p and b change with time in a dynamic environment? Or p might also be a function of f, giving implicit feedback. These estimations might start with a modified Delphi method, taking advantage of diverse experience and expertise within your organization.[11,12]
7. value of perfect valuation information.[11] What if I do not know accurately the probabilities of success or the payoffs? Further, how is profitability impacted if I use the wrong parameters—including due to uncertainty or error—for the KC? What is the value of obtaining better estimates for return and probability? This is in part a sensitivity analysis for the profit. This is related to estimating p and b.

https://doi.org/10.1021/acs.iecr.1c00511
Ind. Eng. Chem. Res. 2021, 60, 7689−7699

8. Size of the investment portfolio. When $\sum f_i > 1$, this can be a call to increase the initial bankroll W_0 for innovation. One can optimize the portfolio of f values, subject to the total anticipated median CAGR being greater than some value chosen by the company. This provides a more rigorous way for establishing the size of an initial innovation amount (W_0), for instance, compared with Knott.[14]
9. Hierarchical use of the KC. If we can use the KC to allocate our research portfolio, can we also use it to evaluate hierarchically? For instance, what if we look up a level for corporate-level decisions (e.g., should we allocate to R&D, better manufacturing, marketing, or M&A) or a lower level about betting on particular questions and hypotheses within a research project? Yes, the KC applies to these also.
10. Criteria other than money. Can a social benefit organization or an untenured professor use the KC method for where to place bets on time allocation? How do we build an appropriate utility function?

The principles in this article can in fact be applied broadly in life, for instance, in analyzing the use of seatbelts in cars, cheating on exams or taxes, or doing fun activities that have a level of danger. Whenever the result depends on a serial set of consecutive and mostly independent events, the KC or a variant is a useful way to examine the problem.

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at https://pubs.acs.org/doi/10.1021/acs.iecr.1c00511.

Additional information on deriving the KC, calculating values using the continuum Gaussian approximation for the binominal distribution, and translating the return and investment into b and p (PDF)

■ AUTHOR INFORMATION

Corresponding Author

Darrell Velegol − *Penn State University, Department of Chemical Engineering, University Park 16802, Pennsylvania, United States; The Knowlecular Process Company, State College 16803, Pennsylvania, United States;* orcid.org/0000-0002-9215-081X; Email: darrell@knowlecular.com, velegol@psu.edu

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.iecr.1c00511

Notes

The author declares the following competing financial interest(s): Darrell Velegol is President of the Knowlecular Processes Company, which does consulting in innovation processes.

■ ACKNOWLEDGMENTS

I thank my colleague Monty Alger, who has discussed with insight many of the issues in this article. I thank Ashish Kulkarni, who raised the concept to me of how the KC applies hierarchically across the entire company. I also thank Abhishek Kar, who provided important insights into the practical application of the KC.

■ REFERENCES

(1) Shannon, C. E. A Mathematical Theory of Communication. *Bell Syst. Tech. J.* **1948**, *27*, 379−423.
(2) Ulwick, A. W. *Jobs to Be Done*; Idea Bite Press, 2016.
(3) Velegol, D. How to Innovate During a Crisis. *Chem. Eng. Prog.* **2020**, 52−56.
(4) Kelly, J. L. A New Interpretation of Information Rate. *Bell Syst. Tech. J.* **1956**, *35*, 917.
(5) Thorp, E. O. Portfolio choice and the Kelly criterion. *Proceedings of the Business and Economics Section of the American Statistical Association*, 1971, pp 215−224.
(6) Video by Nassim Taleb. https://www.youtube.com/watch?v=91IOwS0gf3g&feature=youtu.be. At 6:50, he says, "this destroys tens of thousands of papers in risk aversion." (Accessed 12 May 2021).
(7) MacLean, L. C.; Thorp, E. O.; Ziemba, W. T. *The Kelly Capital Growth Investment Criterion*; World Scientific Publishing: Singapore, 2011. This book contains original articles (see even its Table of Contents) by some of the most prominent names who have worked in this field, including Nobel prize winners. It introduces the field from its beginnings and follows advances up through the early 2000s.
(8) Samuelson, P. A. The "Fallacy" of Maximizing the Geometric Mean in Long Sequences of Investing or Gambling. *Proc. Natl. Acad. Sci. U. S. A.* **1971**, *68*, 2493.
(9) Ziemba, W. T. A Response to Professor Paul A. Samuelson's Objections to Kelly Capital Growth Investing. *J. Portfolio Manag.* **2015**, *42*, 153.
(10) This is like the St. Petersburg paradox.
(11) Hubbard, D. *How to Measure Anything*; Wiley: New York, 2014. Douglas does not use the term "delphi", but his calibration methods for people are similar.
(12) Sutherland, J. *Scrum: The Art of Doing Twice the Work in Half the Time*; Crown Business: New York, 2014.
(13) Miremadi, M.; Musso, C.; Oxgaard, J. *Chemical Innovation: An Investment for the Ages*; McKinsey on Chemicals, May 2013, pp 1−9.
(14) Knott, A. M. The Trillion Dollar R&D Fix. *Harvard Business Review*, May 2012; Vol. *76*. Knott's article uses a modified Cobb−Douglas production function, using not only labor and capital but also research funding.
(15) Matson, J. *Innovate or Die!*; Paradigm Press, 1996.

https://doi.org/10.1021/acs.iecr.1c00511
Ind. Eng. Chem. Res. 2021, 60, 7689−7699