

# A Letter to New Graduates

By Tom Coyne

Published in *The Index Investor*, August 2009

*Many of you are probably planning to take some vacation time this month, possibly to attend reunions where you will catch up with your family's newest graduates. With that in mind, I thought you might find interesting a letter I recently wrote to one, that also holds some important reminders for those of us who are faced with the challenge of making good asset allocations decisions:*

*Dear Graduate,*

While you have no doubt received an excellent education, it has been my experience that formal schooling too often fails to spend enough time on four points that you will find critically important over the next thirty years. With that in mind, I've written this letter to cover them for you.

## Point 1: Thinking Clearly in an Age of Information Overload

Like me, most of your teachers grew up in the age of information scarcity. When we had a paper to write, we trundled off to the library, where we used ancient technologies like the Dewey Decimal System, Card Catalogues, and Encyclopedias to search for whatever information we could find about a subject. Finding this information, and then organizing it (which most of did by taking laborious notes on what we read) usually ate up more than half the time we spent on a project. In many cases, the hypotheses we developed depended on the data we could find. On the other hand, when it came to supporting or refuting these

hypotheses, we became quite good at combining limited data with assumptions and logic to build well constructed arguments.

Today, we live in a different world, where a few mouse clicks bring thousands of pieces of data about a subject to your desktop, where it can be cut and pasted into information organizers. In this environment, the old approach of trying to extract meaning from relatively limited data will quickly be swamped by information overload. Clearly we need a new approach that is based on the assumption of information abundance, rather than information scarcity. In broad terms, there are two paths you can take. The first is essentially a computer assisted version of the old fashioned approach. It goes by many names, including “data mining”, “knowledge discovery”, “exploratory search systems” and “machine learning”, and is based on software programs that can analyze terabytes of data and extract novel insights from them, including hidden relationships and improved predictions. However, while this approach is undoubtedly powerful, its practical application is usually limited by time, cost and limited data availability. This brings me to the second approach: thinking like Karl Popper and Thomas Bayes. Let me explain. Say there are two hypotheses: A and B, and I asked you which was more likely. How would you develop your answer? Karl Popper is famous for his doctrine of falsification – that the essence of the scientific method lies not in proving hypotheses, but in disproving them (there’s even free software available that helps you with this process – see <http://www2.parc.com/istl/projects/ach/ach.html>). Hence one way to approach this issue would be to seek information that would disprove each of the two hypotheses. Bayes would say you need to go further than this, and assess the potential diagnostic value of the information you seek. For example, a piece of evidence that you would see 60% of the time if A were not true, and 50% of the time if B were not true is much less valuable than a piece of evidence that you would expect to see 95% of the time if A were not true, but only 5% of the time if B were not true. More formally, Bayes teaches us to seek evidence with a high “likelihood ratio” – that is, evidence that is much more likely in the case of one hypothesis (e.g., A is not true) than the others being evaluated.

Thinking clearly in the face of information overload also requires us to develop insightful initial hypotheses to guide our search for information. While these will often be refined in an iterative manner as the information search proceeds (for more on this, see “Information Foraging” and “The Sensemaking Process and Leverage Points for Analyst Technology” by Pirolli and Card), starting out with good hypotheses often saves a substantial amount of time, and leads to a better analytical result. So how do you develop good initial hypotheses? In many cases, we have existing mental models that enable us to quickly generate them. These mental models are usually based on a combination of theories and practical experience, and encode our memories of the most important cause and effect relationships, information cues, and potential outcomes of different courses of action.

However, most people’s mental models suffer from some well-known shortcomings: as humans, we struggle to understand situations in which (a) cause and effect are widely separated in time; (b) cause and effect have a non-linear relationship, and/or (c) in which effects have multiple causes that are themselves related to each other. Given this, while an existing mental model will usually provide us with at least one initial hypothesis about a situation, it often helps to develop a second hypothesis by identifying (and changing) one or more assumptions that underlie the first one. In my experience, the right ones to focus on are the linchpins – assumptions that are both critical to the first hypothesis, and highly uncertain. Alternatively, in a competitive situation, where the first hypothesis is an opponent’s most likely course of action, a good alternative hypothesis is the opponent’s most dangerous (to your goals) course of action. Finally, if you are truly confronted with a unique situation where existing mental models don’t apply, three approaches can guide your hypothesis development: analogies, deductions from theories that may apply, and/or reasoning about the players involved (e.g., their goals, constraints, perceptions, and possible courses of action).

## Point 2: Accepting the Non-Trivial Impact of Luck in Life

As human beings, we have a fundamental need to maintain some sense of control over our lives (as uncertainty triggers a primal fear), as well as a positive self-image. Unfortunately, these very human traits result in a distorted understanding of the role of luck (or randomness, if you prefer that term) in our lives. To put it succinctly, when we win, we attribute too much of our success to our own skill, and not enough to good luck. And when we lose, we assign too much of the blame to bad luck, and not enough to our own performance. In both cases, our misperception of the role of luck causes us to spend too little time examining how we need to adapt our thinking and/or behavior. Unfortunately, failure to adapt usually catches up with us, in unpleasant ways. A related problem is our tendency to misperceive random sequences. As Rabin and Vayanos note (in “The Gambler’s and Hot-Hand Fallacies”), the gambler’s fallacy derives from a fallacious belief in the “law of small numbers”: people believe that a small sample should resemble closely the underlying population, and hence believe that heads and tails should balance even in small samples.” An example of this is underreaction to short-term streaks, in the belief that they will soon reverse, failing to recognize that a random process (e.g., coin flipping) will naturally generate such short term streaks, and that the expected statistical outcome (50% heads, 50% tails) will only be apparent after a large number of trials. “On the other hand, people also sometimes predict that random sequences will exhibit persistence rather than reversals – e.g., the belief that a basketball player [or an active investment manager] has a ‘hot hand.’”. In this case, people over-react to long-term streaks, taking them as evidence that the underlying process is not random, even when it is -- i.e., that the observed results reflect skill and not just luck. To summarize: our difficulty with accurately perceiving randomness leads us to overestimate the probability a short-streak will reverse, and over-estimate the probability that a long streak will continue (for more on the role of luck, read these two good books: Fooled By Randomness by Nassim Nicholas Taleb, and The Drunkard’s Walk, by Leonard Mlodinow).

Point 3: Making Good Decisions in the Face of Uncertainty

We'll assume a simple example. You face a situation in which you could follow course of action (or make choice) "A", "B", or "C". The result of each choice (e.g., whether you avoid a crisis, or how much money you make) depends on the conditions that prevail in the future. Let's call these scenarios 1, 2, and 3. You can organize this using a simple table. In the following example, I've assumed a money payoff as the result under each scenario.

	<b>Scenario 1</b>	<b>Scenario 2</b>	<b>Scenario 3</b>
<b>Choice A</b>	\$100	\$(50)	\$50
<b>Choice B</b>	\$50	\$200	\$10
<b>Choice C</b>	\$(60) – you lose \$60	\$130	\$300

The classical theory of decision making in the face of risk shows you how to identify the best choice in this situation. It works like this: (a) Assign a probability to each scenario, so that the probabilities total to 100%. (b) Multiply each result by the probability of the scenario. (c) Add up the resulting numbers for each choice to obtain its expected value (that is, its expected result). (d) Select the choice with the highest expected value. The following table shows an example of this:

	<b>Scenario 1 (25%)</b>	<b>Scenario 2 (50%)</b>	<b>Scenario 3 (25%)</b>	<b>Expected Value</b>
<b>Choice A</b>	\$100 x 25% = \$25.00	\$(50) x 50% = \$(25.00)	\$50 x 25% = \$12.50	\$25.00 - \$25.00 + 12.50 = <b>\$12.50</b>
<b>Choice B</b>	\$50 x 25% = \$12.50	\$100 x 50% = =\$50	\$10 x 25% = \$2.50	\$12.50 + \$50 + \$2.50 <b>= \$65.00</b>
<b>Choice C</b>	\$(60) x 25% = (\$15.00)	\$130 x 50% = \$65.00	\$300 x 25% = \$75.00	= \$15.00 + 35.00 + 75.00 <b>= \$125.00</b>

As you can see, classical decision theory in the face of risk tells you that "C" is the best choice, followed by "B", and then by "A" in last place. However, classical

decision theory has some limitations. The first is that the scenario probabilities are usually subjective. That means that two people who agree on the three choices could still disagree over which one is best simply because they believe the Scenarios have different probabilities.

But an even bigger problem is the way different people interpret the concept of “risk.” For example, rather than using the probabilities of the different scenarios as the measure of risk, another person might think of risk as the size of the range of possible outcomes -- \$150 for Choice A, \$190 for Choice B, and \$360 for Choice C. He or she might then want to choose on the basis of the ratio of the expected value of the choice to the range of possible outcomes (in which case, Choice B and Choice C look about equivalent). Alternatively, suppose two people, say, John and Jane, differ in their willingness to lose money. John chooses “C” because he is willing to risk a loss of \$60 to achieve the maximum expected gain of \$125. However, Jane isn’t willing lose any money, and so chooses “B”, which has a lower expected value of \$65. Classical decision theory might say that Jane made an irrational decision. However, it makes perfectly good sense in light of research that shows how losses typically have twice the emotional impact as gains.

But there is an even bigger problem with classical decision theory, that goes beyond different definitions of “risk”. The problem is this: when we have to make decisions in the real world, more often than not we don’t know all the possible future scenarios, much less their probabilities. In this situation, we have to make decisions in the face of “uncertainty” not “risk.” So how do you approach this?

The first question to ask is “do I have to make this decision now?” In the face of uncertainty, if you can wait to make a decision, it is usually to your advantage, as you will have more time for “sensemaking” – developing a better understanding of the scenario that is developing. Unfortunately, in many cases, you can’t wait, and have to choose. In these situations, and particularly when you don’t have much time to make a decision, the first question you should ask is “do I recognize this situation?” If you do, then that recognition should trigger the memory of what you did when you encountered it before, and how that course of action turned out. If it turned out well, you will probably decide to do the same

thing again. If it didn't turn out well, you'll probably want to make some changes. This process results in your initial plan.

If you don't recognize the situation, you have to quickly gain a rough understanding of it. To do this, ask yourself three questions: (a) what are the most important elements (people, conditions, objects, etc.) in the situation I'm facing? (e.g., we're having a picnic by the lake; there is a boat on the shore; the car is a half-hour hike away; I see fire on a ridgeline about a mile away; and the wind is blowing towards me); (b) What are the most important relationships between these elements? (e.g.: What is the probability the fire will reach us before we can get back to the car? Will the road be crowded with other cars and slow us down? Is the lake big enough that we could get in the boat, go out into the middle of the lake, and escape the fire?); and (c) How the situation is likely to evolve in the near future? (e.g.: do I see or hear any fire engines? Are other people starting to leave? Did some other group already take the boat?). This three step process develops what is called "situation awareness." Good situation awareness enables you to identify good choices or options, and to choose the one that, in your view, appears to have the highest probability of achieving your goal (in this case, escaping the fire).

The next step is to quickly test your plan before you put it into action. The best way to do this is called a "pre-mortem", which is like making up a little story about your course of action. Assume you are in the future, and your plan has failed. Tell a story (in your mind if you are alone, or have your team do this if you are leading a group) about what went wrong. Usually, the key to this story will be an important assumption in your original plan that turns out to be wrong. Then ask yourself what you could have done differently to enable your plan to succeed. Here's a quick example. "My plan to have a picnic by the lake failed because of a surprise rainstorm. My plan to have a nice day out with my family could have succeeded if I had also made reservations at that nice restaurant by the lake, just in case it rained." Making this reservation is called "hedging" – incorporating alternatives in your plan that you can follow if the original plan goes wrong. Usually, a pre-mortem will cause you to consider more hedges. Sometimes it will

cause you to discard your original plan and make a new one. That's what makes it such a powerful technique to use when you are faced with having to make a decision in the face of uncertainty. Once you have done this, mentally think about how your plan, complete with its hedges, will play out. This is called "mental simulation" and sometimes it too can highlight problem areas where you need to think about hedges, or to make other changes in the plan. Once you have finished this step, you've made your decision in the face of uncertainty, and you're ready to implement it – while constantly checking to see if any of your key assumptions aren't matching up with unfolding reality (which, beyond a certain point, should trigger an adaptation to your plan).

One last, but important point: decision making tires you out, and can trigger some unpleasant consequences, including reduced physical stamina, task persistence, cognitive performance, and self control (see "Decision Fatigue Exhausts Self-Regulatory Resources" by Vohs et al). Keep this in mind when you are faced with a sequence of decisions.

#### Point 4: Using Options to Achieve Your Long-Term Goals in an Era of Rapid Change

I realize that for at least the past decade, you have been constantly told that you live in a period of unprecedented change. However, when I look back at my father's life – which began in 1918 and ended in 2008 – I'm not sure this is the case. Every generation probably thinks that it is facing a faster pace of change than anything that has come before. But what is constant over time is that some people seem to cope with this change better than others. What accounts for this? At one level, it may be biological, in that some people have less intense fear reactions to the loss and uncertainty that are an inescapable part of living through a period of rapid change. At a more conscious level, the ability to cope with change involves emotional maturity and stability (i.e., "grace under pressure"), the social ability to build a web of supportive interpersonal relationships, and the cognitive flexibility to question your assumptions, make sense of changing

circumstances, and constantly adapt your mental models. Finally, at the strategic level, many of the people who achieve their long-term goals in spite of rapid and unpredictable change have consciously focused on creating and nurturing a portfolio of options.

A friend of mine once observed that many decisions in life – and particularly career decisions -- have an option element to them. Some decisions create more possibilities than they close off (i.e., they have the profile of buying a financial call option), while others close off more possibilities than they create (i.e., they have the profile of selling a financial put option). In the face of rapid change, seek out the former type, and try to avoid the latter. Diversification is just as important in a career and other aspects of life as it is in an investment portfolio. Of course, the other side of this issue is that you need to retain the flexibility – cognitive, emotional, social and financial – to exercise an attractive option when it is “in the money.” There are lots of different aspects to this, including clarity about your long-term goals, sharing them with the most important people in your life, and resisting the temptation to forego long-term financial flexibility for higher short-term consumption (as my father used to say, “the real secret of feeling rich isn’t having more, it’s wanting less” – but that’s a subject for another letter).

In sum, I have no doubt that over the years of your schooling you have heard many references to “effectiveness” and “efficiency.” Indeed, these are two of the three standards against which the performance of all organisms and organizations can be measured. Unfortunately, the third standard – adaptability – is probably the most important in a period of rapid change, where accurate foresight is next to impossible, particularly as the time horizon lengthens. Yet I’d bet that it is the one you’ve learned the least about. Hopefully, this letter will make up for that, and help you achieve your goals over the next forty years. Last but not least, I wish you much good luck (it helps)!