The Glide Path "Takeoff"

Dimitry Mindlin, ASA, MAAA, Ph.D.
President
CDI Advisors LLC
dmindlin@cdiadvisors.com

April 15, 2015

ABSTRACT

This paper demonstrates that a complete glide path may need an equity "takeoff" – increasing equity allocations initially and decreasing allocations after a "peak." A glide path with a proper "takeoff" can generate better outcomes than a conventional glide path whose equity allocations never increase. These findings defy the broadly accepted rule "more stocks when younger, more bonds when older." The paper advocates a disciplined outcome-driven approach to the design of optimal glide paths.

George Orwell once said, "*We have now sunk to a depth at which restatement of the obvious is the first duty of intelligent men*." The spirit of this quote is fully applicable to the current situation in the area of optimal glide path design.

The concept of *glide path* is increasingly becoming one of the major aspects of portfolio management. A glide path represents a long-term investment strategy that includes the current and future portfolios. Glide paths are vital for the management of DC plans, 529 plans, and other investment programs.

The problem is most glide paths lack a disciplined outcome-driven approach. The theory of optimal glide path design is still in its infancy.[1] Few glide path designers attempt to justify the evolution of portfolios in their glide paths. In particular, the evolution of equity allocation for most glide paths appears to be based on little more than "rules of thumb and folklore."

The lack of a disciplined approach is related to the most essential properties of optimal glide paths. As an example, most glide paths designers follow the rule "more stocks when younger, more bonds when older." Thus, their glide paths' equity allocations never increase. Yet, this rule's foundation is shaky at best. One notable exception to this rule and a couple of attempts to justify it are discussed later in the paper.

The objective of this paper is twofold.

1. To present a case study that demonstrates that a glide path whose equity allocations are increasing initially and decreasing after a "peak" is superior to the conventional glide path whose equity allocations are non-increasing. Moreover, the case study contains the estimates of "surcharges" the conventional glide path imposes on plan participants.
2. To sketch out a disciplined quantitative framework to optimal glide path design and compare it to other approaches.

The criteria for glide path comparison in this paper is based on the following principle that we consider self-evident. *A better glide path generates better outcomes*. We can debate the scientific quality of various frameworks after we analyze the outcomes.

Without further ado, let us proceed to the case study.

A Case Study

The choice of this case study was driven primarily by this author's goal to make it easy to understand and replicable. Another goal was to ensure that all calculations can be performed in a modest size Excel workbook.[2]

John is a 25 year old DC plan participant that currently makes $40,000 annually and has $2,000 in his retirement account. John has made a commitment to contribute 10% of his income until his retirement at 65. His desired replacement ratio to be funded by the DC plan is 40%. John assumes that he will pass away at 95.

John wishes to have a long-term investment strategy – a glide path – that uses just two asset classes: stocks and bonds. He is considering two choices of glide paths. The first one is the conventional glide path presented in *Exhibit 1*.

Exhibit 1

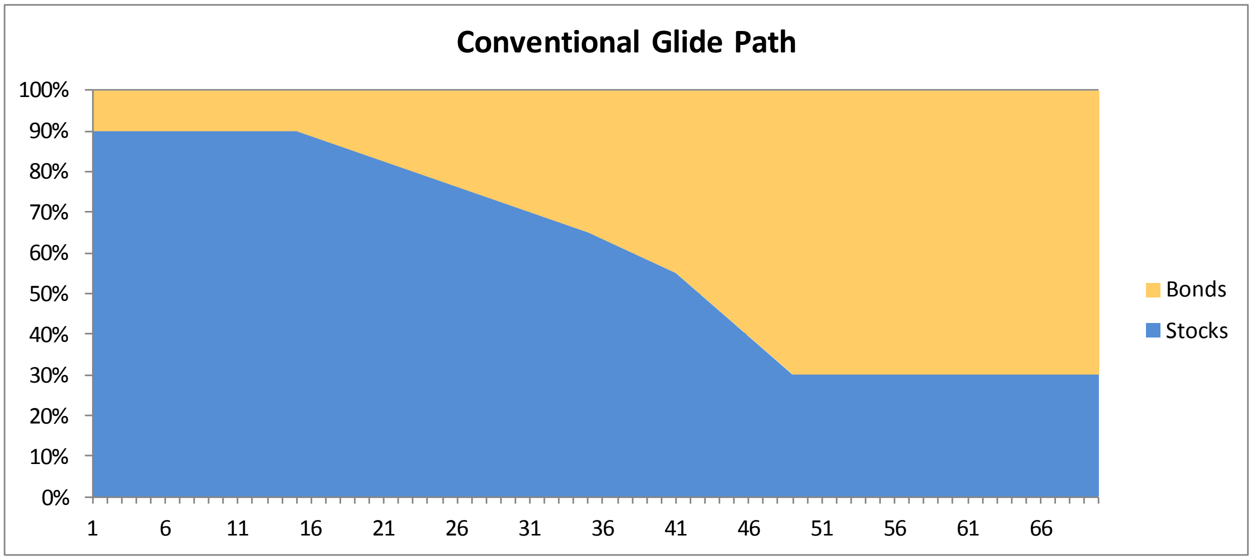

The second one is generated by the glide path optimizer designed according to the principles of *Commitment Driven Investing* (CDI) and presented in *Exhibit 2*. The design methodology is discussed later in the paper.

Exhibit 2

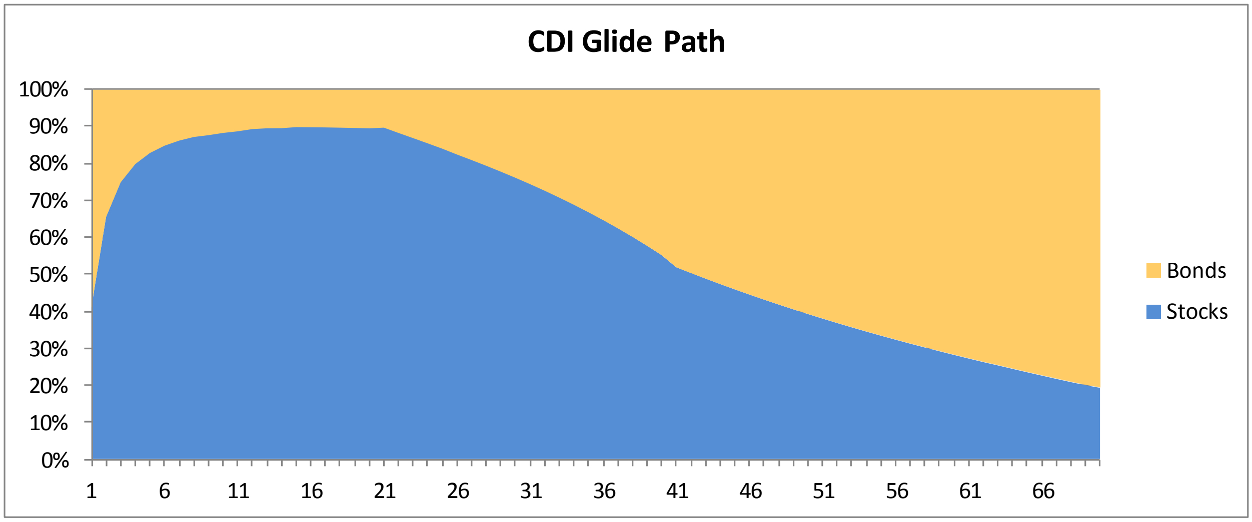

Exhibit 3 compares the evolutions of equity allocations for these glide paths.

Exhibit 3

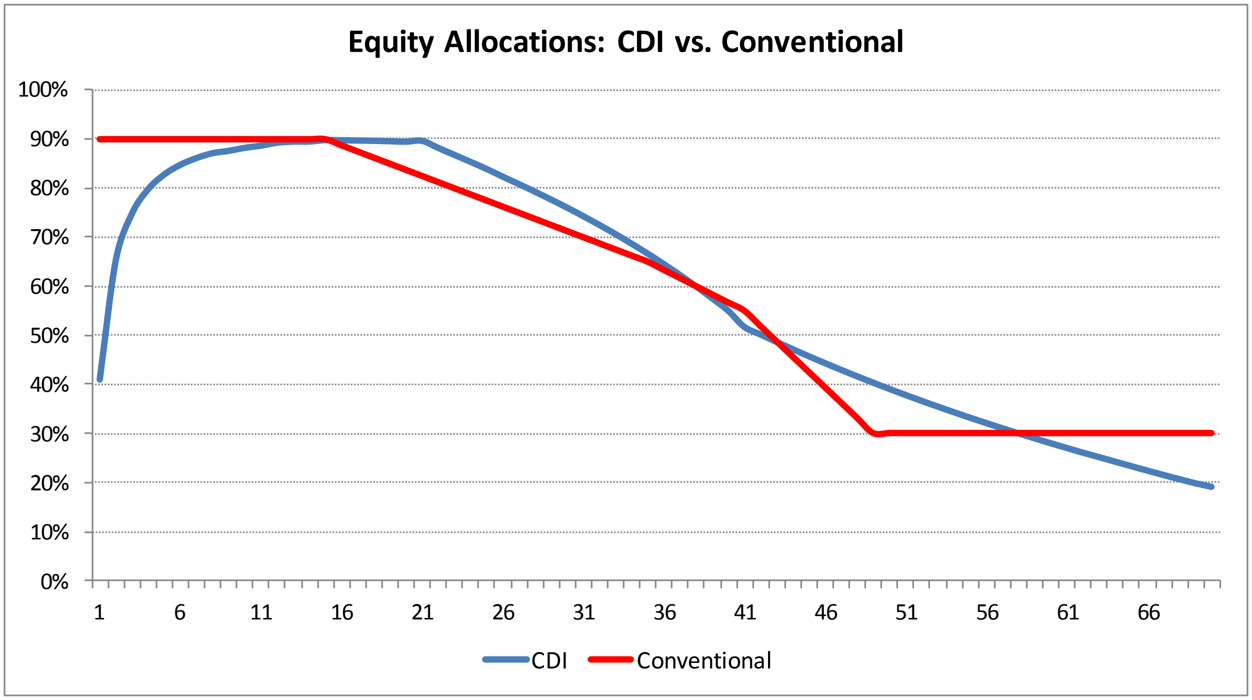

The conventional glide path has an annualized mean return of 6.52%, standard deviation of return of 10.80%, and average equity allocation of 60%. The CDI glide path is specifically designed to have the same annualized mean return, standard deviation and average equity allocation. Therefore, the conventional and CDI glide paths are similar *in the "asset-only" space*. It should be noted that there is a multitude of CDI glide paths with this property; the selected CDI glide path is just an example. Also in this example, intra-year portfolio selections play no role since there are only two asset classes. Any difference in outcomes is due to *different evolutions of equity allocations* in these glide paths.

Yet the conventional and CDI glide paths are quite dissimilar *in the "asset-commitment" space*. In particular, they differ in the timing of higher equity allocations. For example, the conventional glide path has 85% of equities or more in years 1 to 19. A similar time period for the CDI glide is 7 to 24. Also, the CDI glide path has a distinct equity "hump" – the equity allocations increase from year 1 to year 17 and decrease thereafter.

To analyze the outcomes of this funding problem, we focus on shortfall probabilities mostly due to the simplicity of this concept. Other measurements are also important but outside of the scope of this paper.

In this paper, shortfall probabilities and other measurements (e.g. future asset values) are estimated using the methodology of simulation-free stochastic analysis developed at *CDI Advisors*. In contrast, most glide path designers employ Monte-Carlo simulations to estimate these measurements. However, Monte-Carlo simulations possess certain shortcomings that we would prefer to avoid in this paper. The results of the simulation-free stochastic analysis is further validated using Monte-Carlo simulations, see *Appendix 1*. The capital market assumptions are presented in *Appendix 2*.[3]

Let us look at the ability of these glide paths to fund John's desired replacement ratio of 40%. The shortfall probabilities for the conventional and CDI glide paths are 54.1% and 51.9% correspondingly. Thus, the CDI glide path is more likely to fund John's desired standard of living in retirement.

This result, however, is vastly insufficient to pronounce the CDI glide path superior to the conventional one. The requirement for additional considerations lies at the heart of the glide path concept.

We assume that John will rebalance his *future* portfolios according to the glide path selected *today*. Imagine that John is revisiting his asset allocation at some point in the future. John does not have to follow the glide path he selected in the past. He is at liberty to select any glide path he deems optimal.

This observation implies one of the key principles of optimal glide path design. *Every "sub"-glide path of an optimal glide path must be optimal on its own.* Otherwise, the glide path assumption is simply unrealistic. We should expect John to reevaluate his glide path regularly and accept optimal glide paths only.[4]

Therefore, we should compare the shortfall probabilities for the conventional and CDI glide paths throughout John's lifetime. Obviously, we do not know the "starting" asset values to be used with future sub-glide paths, but we can make reasonable assumptions about them.

As a first step, we assume that the starting asset value after N year is equal to its mean value. Under this assumption, we can calculate the shortfall probabilities for all N. The results are presented in *Exhibit 4*.

Exhibit 4

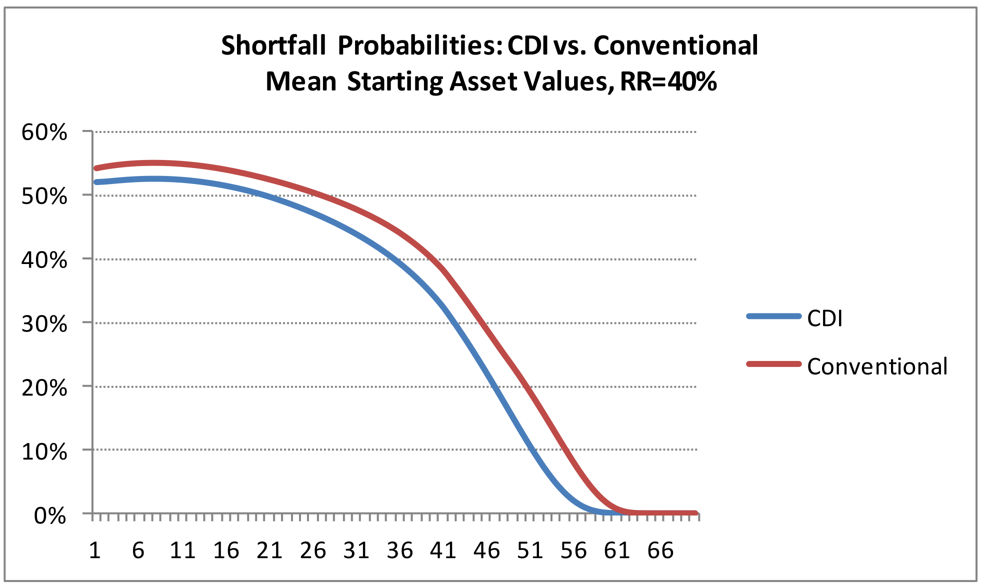

This exhibit demonstrates that the shortfall probabilities generated by the CDI glide path are lower in all years. To show this result from a different angle, *Exhibit 5* shows the magnitude of the differences between these shortfall probabilities.

Exhibit 5

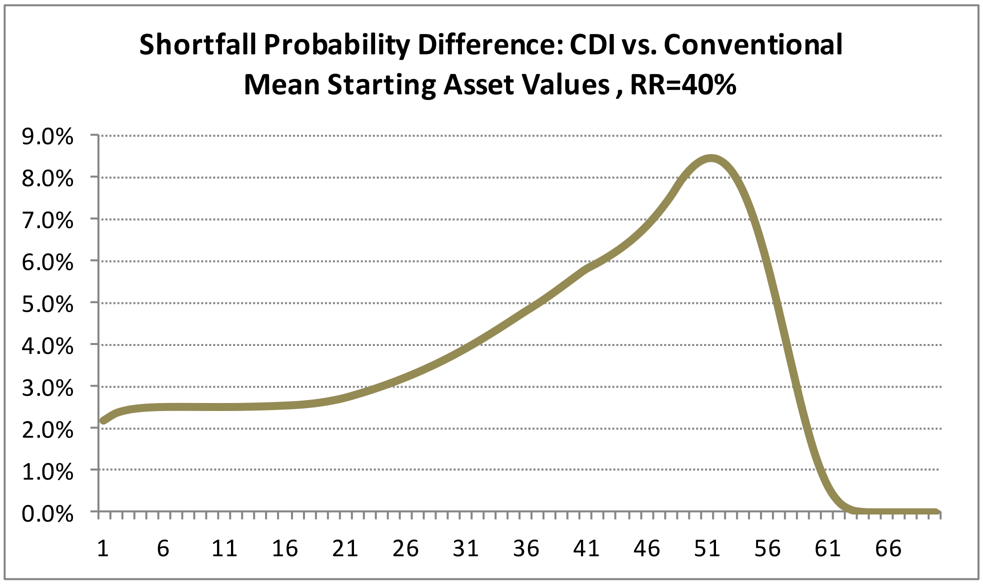

It should be noted that the mean asset values may be overly optimistic in a sense that the probability of achieving them can be lower than 50%. Therefore, we should look at various percentiles of starting asset values that represent a wide range of outcomes. *Exhibit 6* presents shortfall probability differences for the 5th, 25th, 50th, 75th, and 95th percentiles of starting asset values every year. *Exhibit 6* is similar to *Exhibit 5*, but uses the percentiles rather than the means of starting asset values.

Exhibit 6

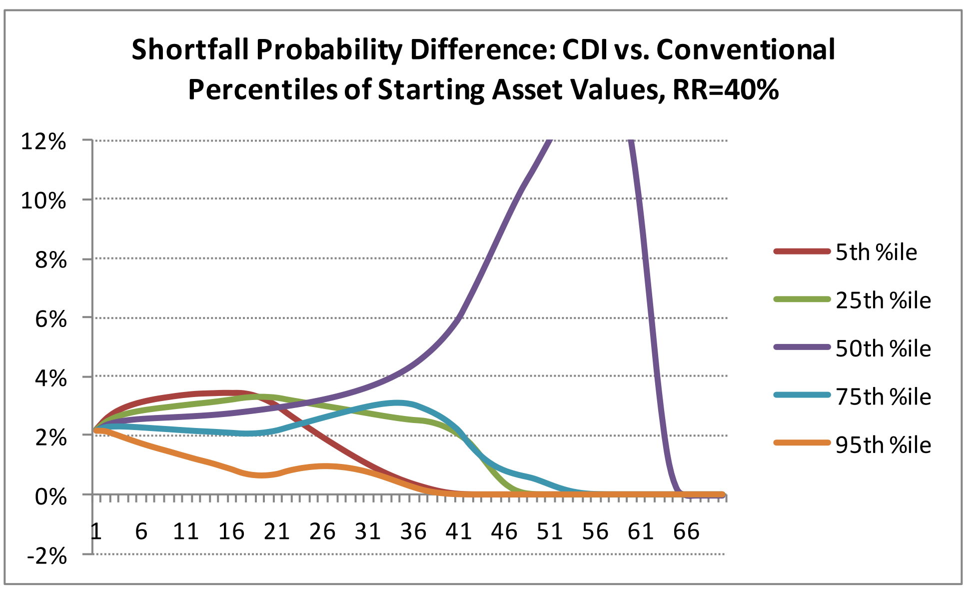

Again, the shortfall probabilities generated by the CDI glide path are lower for a wide range of starting asset values – between the 5th and 95th percentiles – in all years.

Another measurement of outcomes is the distribution of terminal asset values (the end of the glide path period). Selected percentiles of this distribution for both glide paths are presented in *Exhibit 7*.

Exhibit 7

Terminal Asset Values						
RR=40%						
	Mean	5th %ile	25th %ile	50th %ile	75th %ile	95th %ile
CDI	973	-3,175	-1,876	-263	2,355	9,230
Conventional	662	-3,223	-1,974	-451	1,984	8,251

As we see, the selected measurements of the terminal asset value distribution (the mean and the percentiles from 5th to 95th) generated by the CDI glide path are generally higher than the corresponding values generated by the conventional glide path.

This discussion leads the following important questions. What are the implications of glide path selection for plan participants? What is the impact on contribution rates, replacement ratios, and required returns?

To answer these questions, we need an outcome measurement for the conventional glide path that would be equal to its CDI counterpart with adjusted contributions, replacement ratio, or returns. As a starting point, let us use *the median terminal asset value* for this purpose.

We calculate the increase in contribution rates, the reduction in replacement ratio, and the increase in the required return with the conventional glide path that would produce the same median terminal asset value for the unadjusted values with the CDI glide path *ceteris paribus*. The results are presented in *Exhibit 8*.

Exhibit 8

Terminal Value Median Matching			
Adjustments to Be Used with Conventional Glide Path			
RR=40%			
Terminal Value Median	Contribution Rates	Replacement Ratio	Additional Return (bps)
-451	0.0%	0.0%	0
-263	0.3%	0.0%	0
-263	0.0%	-1.2%	0
-263	0.0%	0.0%	10

Here is how to read this exhibit. The conventional glide path generates the median terminal value of -451. We analyze the following three ways to bring it up to -263 (the median terminal value generated by the CDI glide path): increasing the contribution rates, lowering the replacement ratio, increasing the mean expected return for each asset class. To get the desired median terminal value, we could increase the contribution rates by 0.3%, or reduce the replacement ration by 1.2%, or require additional return of 10 basis points for each asset class. The last adjustment may also be interpreted as the required reduction of the asset management fee.

It should be noted that the median terminal value does not directly incorporate the substantial upside generated by the CDI glide path. In contrast, the mean of the terminal value does incorporate the upside. Since different choices of "matching" measurements may produce different answers, it is informative to see similar adjustments generated by the mean terminal value as the matching measurement. The results are presented in *Exhibit 9*.

Exhibit 9

Terminal Value Mean Matching			
Adjustments to Be Used with Conventional Glide Path			
RR=40%			
Terminal Value Mean	Contribution Rates	Replacement Ratio	Additional Return (bps)
662	0.0%	0.0%	0
973	0.5%	0.0%	0
973	0.0%	-2.0%	0
973	0.0%	0.0%	12

Here is how to read this exhibit. The conventional glide path generates the mean terminal value of 662. We analyze the following three ways to bring it up to 973 (the mean terminal value generated by the CDI glide path): increasing the contribution rates, lowering the replacement ratio, increasing the mean expected return for each asset class. To get the desired mean terminal value, we could increase the contribution rates by 0.5%, or reduce the replacement ration by 2.0%, or require additional return of 12 basis points for each asset class. All of these adjustments ("surcharges") are meaningful.

It should be emphasized that this case study was specifically designed to give the conventional glide path a fighting chance. The use of just two asset classes renders optimal portfolio selection – one of CDI's strongest suits – underutilized. The matching annualized portfolio risk and return for the conventional and CDI glide paths ensure that the comparison is "apples-to-apples."

Yet even under these constraints the CDI glide path generates better outcomes. Specifically, the CDI glide path is aggressive when it matters – when the accumulated assets are large enough to take advantage of higher expected returns. The CDI glide path is also conservative when it matters – in the "takeoff" phase where the assets are in the process of initial accumulation.

Without these self-imposed constraints – when the glide path designer is at liberty to add additional asset classes that may be beneficial to plan participants and vary the shape of the glide path to generate the best outcomes – magnitude of the "surcharges" imposed by the conventional glide path should be much higher. These issues will be the subject of future research.

Let us recap these results. The shortfall probabilities generated by the CDI glide path are lower than the corresponding probabilities for the conventional glide path for the broad ranges of starting asset values in all years. The selected measurements of terminal asset values generated by the CDI glide path are higher than the corresponding values for the conventional glide path. The inefficiencies of the conventional glide path impose meaningful "surcharges" on the plan participants that utilize the glide path. Each of the measurements analyzed in this section is inconclusive by itself. In aggregate, however, this section makes a compelling case that the CDI glide path is superior to the conventional one in the "asset-commitment" space.

CDI Glide Path Design

While the previous section demonstrates that the CDI glide path performs better than the conventional one, it says little about the reasons for the better performance. The primary reason – a superior glide path's design methodology – is presented in this section. This "order of operations" – outcome analysis first, methodology analysis next – is used to highlight the importance of the outcome analysis.

As mentioned before, the CDI glide path is generated by the optimizer that utilizes the principles of CDI. This section does not attempt to present CDI in full; it presents a brief review of the key concepts of CDI related to the case study. For more information, see Mindlin [2014] (the essentials of CDI), Mindlin [2013B] and Mindlin [2013C] (the principles of optimal glide path design).

One of the key aspects that distinguishes CDI from other frameworks is its emphasis on investment objectives. For DC plans, there are three major types of investment objectives that represent the first step in the development of CDI:

A. To maximize post-retirement spending given contributions and risk.
B. To minimize contributions given post-retirement spending and risk.
C. To minimize risk given contributions and post-retirement spending.

Subsequently, CDI quantifies these objectives and ascertains relationships between them. The choice of the right objective depends on the nature of the funding problem.[5]

For the case study, the selected primary objective is to maximize the probability of success given the contribution rate and replacement ratio (objective *C*). While we could attempt to maximize this probability directly, this attempt would be insufficient because it is overly simplistic. There are additional considerations and constraints that must be taken into account.

First, we recognize that the investor will re-examine his asset allocation regularly. Therefore, any glide path the investor selects represents a number of portfolio selections made by the investor's virtual ageing "clones." This setup represents a classic strategic game: we have "players" (the investor's ageing "clones"), their actions (portfolio selections) and preferences (the higher the probability of success, the better). Under common rationality assumptions, all portfolio selections should form a *Nash equilibrium* (NE) strategy.[6]

Second, the starting asset value in a particular year is generated by the starting segment of the glide path; the shortfall probability is heavily influenced by the ending segment. Consequently, the starting and ending segments of a glide path should collaborate effectively in order to produce lower shortfall probabilities in all years.

Third, we would like to have higher probabilities of success for wide ranges of starting asset values in all years. Consequently, the objective function should encompass the full range of outcomes.

Having taken into account all these considerations, *the stochastic present value (SPV) mean-variance* version of CDI has been selected for the case study. This particular version of CDI was inspired by Modern Portfolio Theory (MPT), but differs from MPT in several aspects. Mindlin [2014] has a detailed discussion regarding the similarities and dissimilarities between MPT and the SPV mean-variance version of CDI.

CDI can be considered as an expansion of MPT for investors with financial commitments. However, there is a directional difference between MPT and CDI. Namely, MPT optimizes stochastic *future* values, which are preferred to be *high*. In contrast, this version of CDI optimizes stochastic *present* values, which are preferred to be *low*. In this sense, CDI is MPT “in reverse.”

In a nutshell, the glide path optimization process used to design the CDI glide path is as follows.

1. For a given year N, define RA_N as the SPV of the remaining cash flows from N to the end of the mortality table.
2. The "risk-adjusted expected cost" is defined as the mean of RA_N plus the standard deviation of RA_N times a risk aversion factor.
3. Select the values of the risk aversion factor in all years, see *Exhibit 10*.
4. Using the process of "*backward induction*," minimize the "risk-adjusted expected cost" for a given risk aversion factor in all years. Backward induction optimization starts in the last year and continues backwards to the first year. It can be shown that this process generates NE strategies.

By design, the glide path produced by backward induction optimization generates optimal outcomes in all years. See Mindlin [2013B] for more details.

Exhibit 10

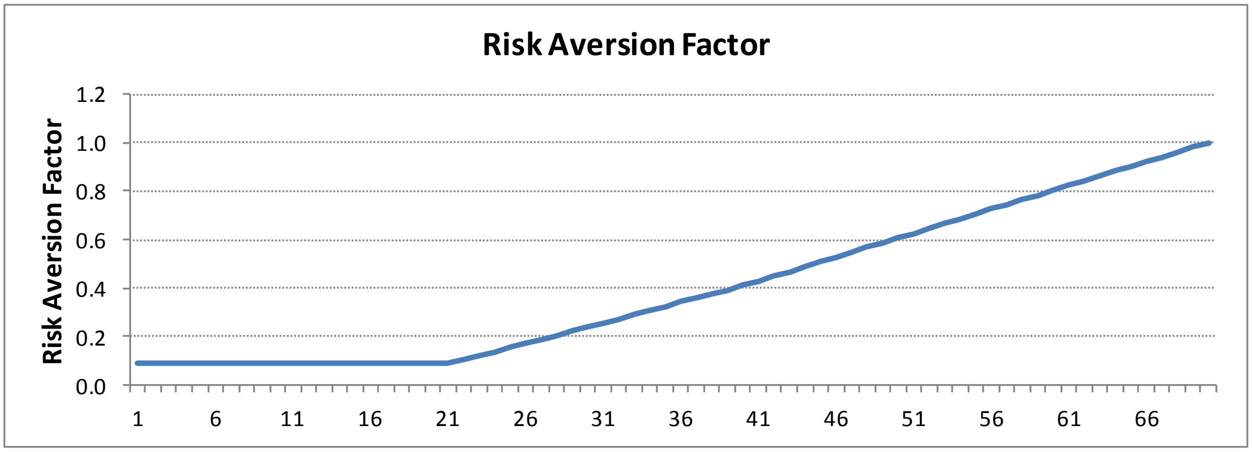

One crucial property of this glide path optimization process may be somewhat concealed but deserves special attention. As emphasized above, the first step in the process is to spell out the investment objective. The type of present value consistent with the objective – namely, the stochastic present value that is based on actual returns – is selected next. This "order of operations" – investment objectives first, present values next – is one of the central elements of this process. This "order of operations" is in fact one the most important aspects that distinguishes CDI from the "human capital theory" discussed later in the paper.

Overall, the ultimate test of any glide path design framework must involve the outcomes it generates. The strongest argument in favor of the CDI based glide path optimization process is resulting glide paths generate better outcomes.

"More Stocks When Younger, More Bonds When Older"?

"Folk wisdom and casual introspection" support the principle "more stocks when younger, more bonds when older".[7] Economists have endeavored to find a proper theoretical foundation for this principle for decades. Yet the case study presented in the previous section demonstrates that "more stocks when younger, more bonds when older" is invalid as a general principle.

This conclusion represents a remarkable departure from the prevailing practices in the industry. The overwhelming majority of TDFs have non-increasing equity allocations in their glide paths. These TDFs may disagree on many aspects of glide path design, but one particular rule enjoys wide-ranging institutional support: equity allocations should never increase over time.[8]

This author is aware of just one notable exception to this rule. National Employment Savings Trust (NEST), a workplace DC pension scheme in the U.K., offers a glide path that effectively has a "takeoff" phase. Even though the NEST glide path methodology appears to have ample room for improvement, NEST may be on to something.[9]

The next section attempts to explore why most glide path designers prefer non-increasing equity allocations.

"Rules of Thumb and Folklore"

Peter Bernstein once observed:

> "*Before Harry Markowitz's 1952 essay on portfolio selection, there was no genuine theory of portfolio construction – there were just **rules of thumb and folklore***."[10] (Emphasis added)

The current situation in the area of optimal glide path design is quite similar to the situation in the area of portfolio design prior to the publication of Harry Markowitz's seminal paper in 1952. The matters discussed in most publications on the subject of glide path design (e.g. asset class selection, active vs. passive investing, fees) are conventional portfolio management issues that are not directly related to optimal portfolio evolution throughout the investor's lifetime.

The design of most TDF glide paths contains plenty of "rules of thumb and folklore." Poorly defined terminology generates abundant confusion. Conflicting "rules of thumb" ignite raging debates sometimes related to matters of questionable meaning and significance (e.g. "to" vs. "through" glide paths). Yet, some of the key aspects of glide path design remain undeveloped.

Specifically, let us consider the basic question of optimal glide path design: *Should optimal glide paths be evolving or stationary?* On one hand, a popular theme in financial folklore compels "more stocks when younger, more bonds when older." On the other hand, as demonstrated in Samuelson [1969], investors under certain conditions should "hold the same fraction of portfolio in equities early and late in life." Accordingly, this question have been contentious for decades.

This section briefly discusses two well-known attempts to answer this question.

The first attempt is a claim that the young should hold more equities because they have more time to recover from adverse market conditions. The "rule of thumb" here is the following: more time means greater risk-taking capacity. Essentially, this "rule of thumb" implies that the riskiness of equities diminishes over time. This alleged property of equities is often called "time diversification."

If the time diversification property of equities were proven, it would conclusively justify the rule "more stocks when younger, more bonds when older," which in turn would justify evolving glide paths. Therefore, the time diversification property is a sufficient condition to justify evolving glide paths with non-increasing equity allocation over time. Do today's glide path designers believe in this property?

In fact, some may.[11] But, for better or worse, the "time diversification" of equities fails as a sound principle of finance. Even a cursory examination of the rule "more time means greater risk-taking capacity" would promptly reveal its superficiality. In fact, "time diversification" of equities is one of the greatest controversies in finance. After numerous publications and debates, the subject remains controversial to this day. A report from *Vanguard* provides a fair assessment of this controversy:

> "*Some of the finest investment minds have participated on both sides, without providing a conclusion. … The debate over time diversification has been long-running and remains unresolved.*"[12]

It should be noted that glide path designers may have no position on the issue of time diversification and still believe in "more stocks when younger, more bonds when older." That is where we have the second attempt to answer the question regarding the optimality of evolving glide paths.

This approach distinguishes two types of assets: financial capital and "human capital." The latter is generally defined as the ability to generate income. To emphasize the role of human capital, this approach is often called the "Human Capital Theory" (HCT).

In a nutshell, HCT treats human capital as a low risk bond-like conventional asset for most investors. For a typical young investor, the value of this bond-like asset is high. To have a well-diversified portfolio, the young investor should have more stocks in his financial capital. As the investor ages, his bond-like human capital decreases and his financial capital increases and shifts toward bonds. According to HCT, this is the main reason the optimal glide path evolves from "more stocks" to "more bonds" for a typical investor.

HCT deserves recognition as arguably the first relatively adequate approach that justifies evolving glide paths. While a detailed analysis of HCT is outside of the scope of this paper, this section offers just brief comments required to compare HCT and CDI.[13]

Human capital is an "asset" in a sense that more of it is better than less of it. This "asset" is a non-tradable non-transferable contingent cash in-flow of uncertain timing and magnitude. The process of optimal glide path design in HCT is based on the assumption that this "asset" can be treated as a conventional financial asset and incorporated into a mean-variance optimizer along with conventional assets.

Thus the first step in the process is to "price" human capital, i.e. to assign a deterministic present value to it. The next step is to define the primary objective, which is to maximize the expected utility of the sum of financial capital and human capital. The optimal glide path is generated via maximization of the expected utility throughout the investor's lifetime.

The shape of resulting glide path closely resembles the shape of the conventional glide path used in the case study. In fact, the intent of the case study was to evaluate the glide path produced by a major provider of TDFs that utilizes HCT. As demonstrated in the case study, the glide path designed according to HCT is sub-optimal. In light of this observation, it is informative to ponder the following question: what accounts for this sub-optimality in the construction of HCT? Which aspects of HCT are responsible for certain inefficiencies in this theory?

To answer these questions, we have identified the following aspects of HCT that look theoretically suspect and appear to belong to financial folklore rather than the science of finance.

First, HCT does not attempt to achieve optimal outcomes directly. Instead, HCT's objective is to maximize expected utility throughout the lifecycle. Outside of financial folklore, it is hard to find any connection between maximizing expected utility and generating optimal outcomes of retirement programs. Moreover, maximization of expected utility as a general objective is somewhat controversial, even though it is widely used in finance and other fields.[14] Some utility functions should work better than others. Utility functions that are not directly related to optimal outcomes of retirement programs may be especially problematic.

Second, the assumption that human capital must have a deterministic "price" is little more than a "rule of thumb." Moreover, we must use a deterministic discount rate to "price" human capital. The very presence of this discount rate creates an internal inconsistency in HCT. Namely, this artificial risk-free "return" on human capital – the discount rate – is used for optimization purposes, but the outcomes are calculated using actual returns. In CDI, human capital is treated as a cash flow, so pricing assumptions are unnecessary. All future and present value calculations consistently use actual returns.

Third, HCT defines this present value before it defines investment objectives. This present value – the price of human capital – is defined first, the investment objective consistent with the present value is defined next. The corresponding "order of operations" in CDI is exactly the opposite. The investment objective is defined first, the present value consistent with the investment objective – the stochastic present value of all financial commitments – is defined next.

Overall, HCT may represent a passable justification of the optimality of evolving glide paths, possess computational conveniences and mathematical elegance. These commendable qualities may explain HCT's popularity among some glide path designers. These qualities, however, should be of secondary importance if this theory generates sub-optimal outcomes.

Conclusion

Discussing the development of the calculus of variations, Paul Samuelson once observed:

> "*All this is good enough for the brilliant eighteenth century. But by the nineteenth it was a scandal that a rigorous mathematical theory was still not known.*"[15]

A scandal, no less. Paul Samuelson had spent over four decades trying to develop a solid theoretical justification for evolving glide paths. He may have felt that "it was a scandal that a rigorous mathematical theory" of optimal glide path design was still undeveloped.[16]

That is where we are today. Too often, glide path designers rely on the "rules of thumb and folklore." "Folk wisdom and casual introspection" are in demand as well. By doing so, certain major components of a complete glide path may be overlooked, namely the "takeoff" phase and future reevaluations.

This paper provides several illustrations to this observation. For a young retirement investor, the conventional glide path not only puts the desired standard of living in retirement at a greater risk throughout the investor's lifetime. Without a glide path "takeoff," it also needlessly exposes the investor to additional equity risk today. A glide path with a proper "takeoff" generated by a disciplined outcome-driven approach may work better for the investor.

This situation should not be characterized as "a scandal" in the spirit of the aforementioned Paul Samuelson's quote. Not yet, at least. However, it may be headed down that path if no changes are made.

What kind of changes would be helpful to the industry? Upgrade the paradigm to a higher standard. Spell out the principles. Identify the objectives. Clarify the language. Emphasize the outcomes.

As George Orwell once noted: "*One's got to change the system, or one changes nothing.*"

APPENDIX 1: Simulation-Free Analysis vs. Monte-Carlo Simulations

In this paper, shortfall probabilities and other measurements are estimated using the methodology of simulation-free stochastic analysis developed at *CDI*

Advisors. The key steps in the methodology include the calculations of moments and selecting appropriate theoretical distributions that match the first three moments of the random variables of interest. These calculations are further validated using Monte-Carlo simulations, see below.

Simulation-free estimates for RR=40% (Exhibit 6):

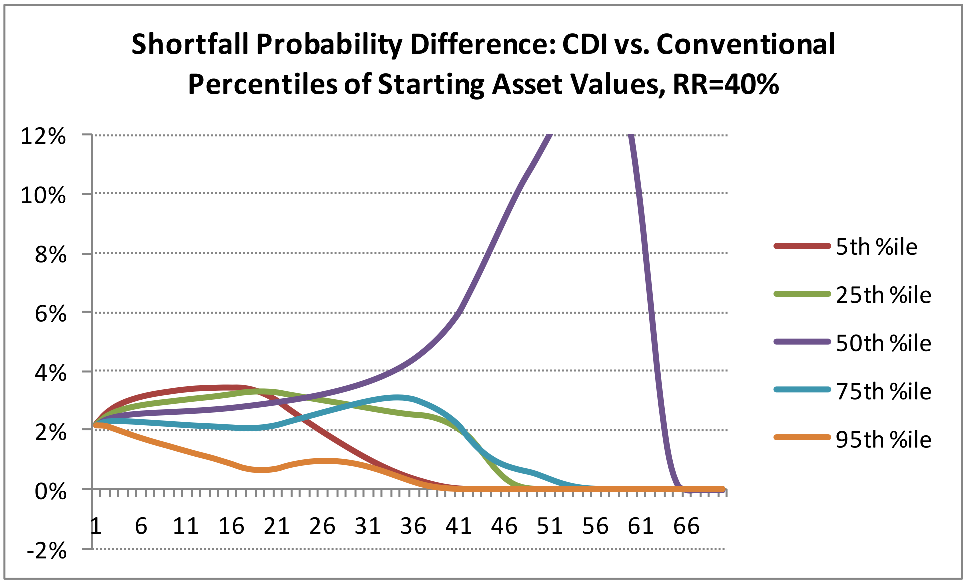

Monte-Carlo simulations based estimates for RR=40%:

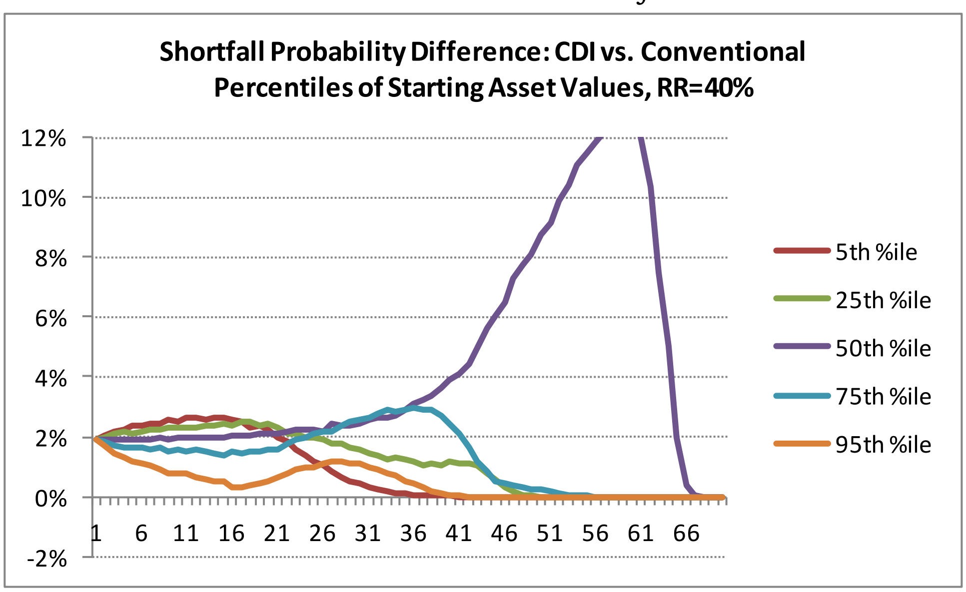

The simulation sample size is 100,000. The results of these simulations corroborate the results of the simulation-free analysis and demonstrate the quality of the approximations utilized in the paper.

APPENDIX 2: Capital Market Assumptions

Return/Risk

	Geometric Mean (%)	Arithmetic Mean (%)	Standard Deviation (%)
Stocks	7.00	8.03	16.00
Bonds	4.00	4.12	5.00
CPI	2.50	2.505	1.00
Wage Growth	3.50	3.505	1.00

Correlation Matrix

	Stocks	Bonds	CPI
Bonds	0.2	1	
CPI	0.0	-0.1	1
Wage Growth	0.1	-0.1	0.7

Portfolio returns are assumed to have lognormal distributions.

REFERENCES

Arnott, R. D., Sherrerd, K. F., Wu, L. J., [2013]. The Glidepath Illusion… and Potential Solutions, *Journal of Retirement*, Vol. 1, No. 2 (Fall) 2013, 13–28.
Bennyhoff, D.G., [2008]. Time Diversification and Horizon-Based Asset Allocations, Vanguard Investment Counseling & Research, *The Vanguard Group, Inc*., 2008, seen on 9/23/2014 here.
Bernstein, P. [2007]. Capital Ideas Evolving, *John Willey & Sons*, 2007.
Bodie, Z., Merton, R., Samuelson, W., [1992]. Labor Supply Flexibility and Portfolio Choice In a Life Cycle Model, *Journal of Economic Dynamics and Control*, 16, 1992.
Donaldson, S., Kinniry F., Aliaga-Díaz, R., Patterson, A., DiJoseph, M., [2013]. Vanguard's Approach to Target-Date Funds, *The Vanguard Group*, December, 2013, seen on 4/14/2015 here.
Estrada, J., [2014]. The Glidepath Illusion: An International Perspective, *The Journal of Portfolio Management*, Summer 2014, 52-64.
Fullmer, R., [2014]. Reflections on Recent Target Date Glide-Path Research, *T. Rowe Price Asset Allocation Insights*, February, seen on 4/14/2015 here.
Ibbotson, R., Milevsky, M., Chen, P., Zhu, K., [2007]. Lifetime Financial Advice: Human Capital, Asset Allocation, and Insurance, *The Research Foundation of CFA Institute*, 2007.
Idzorek, T., [2008]. Lifetime Asset Allocations: Methodologies for Target Maturity Funds, Ibbotson Associates Research Report, *Ibbotson Associates, Inc*., February, 2008, seen on 9/23/2014 here.
Mindlin, D., [2009]. Commitment Driven Investing and Time Diversification, *CDI Advisors Research*, CDI Advisors LLC, 2009, seen on 4/14/2015 here.

Mindlin, D., [2011]. LDI vs. CDI: Bitter Foes or Best Friends? *CDI Advisors Research*, CDI Advisors LLC, 2011, seen on 4/14/2015 here.
Mindlin, D., [2013A]. A Tale of Three Epiphanies, *CDI Advisors Research*, CDI Advisors LLC, March, 2013, seen on 4/14/2015 here.
Mindlin, D., [2013B]. Principles of Optimal Glide Path Design, *CDI Advisors Research*, CDI Advisors LLC, March, 2013, seen on 4/14/2015 here.
Mindlin, D., [2013C]. Glide Path 2.0: The Next Generation, *CDI Advisors Research*, CDI Advisors LLC, April, 2013, seen on 4/14/2015 here.
Mindlin, D., [2014]. Commitment Driven Investing: The Essentials, *CDI Advisors Research*, CDI Advisors LLC, 2014, seen on 4/14/2015 here.
Mindlin, D., [2015]. The Glide Path Assumption, *CDI Advisors Research*, CDI Advisors LLC, 2014, seen on 4/14/2015 here.
Rabin, M., Thaler, R., [2001]. Anomalies: Risk Aversion, *Journal of Economic Perspectives*, 15, 219-32, 2001.
Samuelson, P., [1969]. Lifetime portfolio selection by dynamic stochastic programming, *The Review of Economics and Statistics*, 51, 239-246, 1969.
Samuelson, P., [1970]. What Makes for a Beautiful Problem in Science? *Journal of Political Economy*, December, 1970.
Samuelson, P., [1973]. Proof That Properly Discounted Present Values of Assets Vibrate Randomly, *The Bell Journal of Economics and Management Science*, 369-374, Autumn 1973.
Samuelson, P., [1989]. A Case At Last for Age-Phased Reduction in Equity, *Proc. Natl. Acad. Sci. USA*, Vol. 86, 9048-9051, 1989.
Samuelson, P., [1991]. Long Run Risk Tolerance When Equity Return Are Mean Regressing: Pseudoparadoxes and Vindication of "Businessman's Risk", In Money, Macroeconomics, and Economic Policy: Essays in Honor of James Tobin, Edited by W. C. Brainard, W. D. Nordhaus, H. W. Watts, *MIT Press*, 1991.
Samuelson, P. [1992]. At Last, a Rational Case for Long-Horizon Risk Tolerance and for Asset-Allocation Timing? Active Asset Allocation, edited by Arnott, R.D. and Fabozzi, F.J., *Probus Publishing Company*, 1992.
Scott, B., Donaldson, S., Brancato, M., [2014]. Target-Date Funds: Looking beyond the Glide Path in 2014, *The Vanguard Group*, March 2014, seen on 4/14/2015 here.

ENDNOTES

[1] Idzorek [2008] summarizes the state of affairs in this area as follows: "*... most target maturity equity glide paths lack theoretical substance. ... Little rigorous work has been done to answer how and why the equity-bond glide path should evolve throughout an investor's lifetime and even less work has been done to answer how and why intra-stock and intra-bond splits should evolve over time*."
[2] This worksheet is available to interested parties upon certain conditions. For more detail, contact the author at *dmindlin@cdiadvisors.com*.
[3] As a technical matter, this author is not aware of a closed form solution for the calculation of shortfall probabilities and other measurements.
[4] See Mindlin [2015] for more details.
[5] See Mindlin [2014] for more details.

---

[6] Mindlin [2009] and Mindlin [2013A] contain simplified numerical examples of NE glide paths.

[7] The phrase "folk wisdom and casual introspection" is quoted from Samuelson [1989]. Here is the full quote: "*Maximizing expected utility over a lifetime leads one who has constant relative risk aversion and faces random-walk securities returns to be "myopic" and hold the same fraction of portfolio in equities early and late in life – a defiance of* ***folk wisdom and casual introspection***." (Emphasis added)

[8] While the vast majority of glide path designers advocate downward-sloping glide paths, there have been attempts to promote upward-sloping glide paths, e.g. see Arnott [2013] and Estrada [2014]. Fullmer [2014] appropriately exposes certain fundamental flaws in Arnott [2013] and Estrada [2014], although Fullmer [2014] offers little quantitative support for downward-sloping glide paths.

[9] NEST distinguishes three phases of a participant's lifecycle: foundation, growth, and consolidation. The portfolios in the foundation and consolidation phases have somewhat lower volatility than the ones in the growth phase. As a result, this framework has a glide path "takeoff."

[10] Emphasis added. See Bernstein [2007], page *xii*.

[11] See Scott et al. [2014]: "*Younger investors are more capable of bearing the additional volatility experienced and expected in the equity markets* ..." See Donaldson et al. [2013]:"... *younger investors are better able to take risk than older investors*."

[12] See Bennyhoff [2008].

[13] The origins of HCT can be found in Bodie-Merton-Samuelson [1992]; a comprehensive presentation of HCT can be found in Ibbotson-Milevsky-Chen-Zhu [2007].

[14] Rabin-Thaler [2001] calls the maximization of expected utility "plainly wrong and frequently misleading." The language in Rabin-Thaler [2001] is simply merciless: "... *it is time for economists to recognize that expected utility is an ex-hypothesis, so that we can concentrate our energies on the important task of developing better descriptive models of choice under uncertainty*."

[15] See Samuelson [1970].

[16] See Mindlin [2013A] for more details.