



Rethinking Retirement “Gap” Analysis

Chuck Yanikoski

RetirementWORKS, Inc.

June 2014

Executive summary

Traditional gap analysis for retirement . comparing what you need for retirement against what you’ve already provided for, and then calculating what’s needed to fill the gap between the two . is both commonplace and, very often, counter-productive.

It’s counter-productive because for many people, especially those in the second half of their careers, the gap is so large that filling it looks impossible, not even worth aiming for.

There are two ways to alleviate this problem. The first, which is to fudge the assumptions used in the calculations, is increasingly ineffective as clients get older, because there is insufficient time to fill the gap. The second method, which enables the impact to be spread out over more time, solves the underlying problem by changing the focus in a way that is both effective and much more justifiable.

Although there are dozens, perhaps hundreds, of retirement gap calculators available, few of them use either of these methods, and almost none of them use both. As a result, consumers who actually could solve their retirement problem are discouraged from doing so, and financial companies and their representatives are letting business escape that is already in the palms of their hands . to their own detriment, and that of their clients.

Two potential strategies are presented for improved analytical tools.

Gap analysis, and why it fails.

Traditional retirement gap analysis calculates how much a person or couple should be saving for retirement. Virtually every financial company offers software to perform this calculation, as do other consumer-oriented websites. Manual calculations are described in many books about retirement.

But most of these tools fail because they do not well serve the true purpose of the calculations. In this paper, I analyze why this is so, and what can be done about it.

The *ostensible* purpose of retirement gap analysis is to help people determine how much money they need to save to make their retirement financially secure. If we're honest, though, we'd agree that the *real* purpose involves the profit motive: encouraging people to put more of their earned income into financial products rather than spending it. But this is actually a laudable goal in most cases. Successful efforts to encourage retirement savings benefit consumers by putting them in a better position both to enjoy retirement and to confront the very real financial risks that retirees face, and therefore they also benefit society by reducing the number of indigent elderly citizens who require governmental support. And yes, it also benefits financial companies and their field representatives, but there is no particular harm in that.

The harm arises from *the conflict between the ostensible and the real goals* of the calculations. There are two contributing factors to this opposition. First, sponsors and developers of calculation tools want to keep them simple, so ordinary people are both able and willing to use them (and so that they are not too time-consuming for professional advisors to use). Second, the calculators tend to focus on getting the right answer. Both of these intentions are innocent and even praiseworthy in themselves, but together they tend to produce an unfortunate result. The user of the software does get a number, presented as the correct answer, which the consumer is expected to embrace, but that result may be either motivational (% you increase your savings rate by just 20% over what it is now, you'd be in great shape for retirement!+) or completely de-motivating (% you save 75% of your earned income between now and retirement, you'd be in great shape!+). In the latter case, the typical response is for the consumer to regard the problem as insoluble, and walk away . thus completely defeating the true purpose of the analysis.



Can't gap analysis serve both purposes, though . providing solid answers and also motivating consumers to take action? To some extent, yes, but before setting that as our goal, we need to understand that the attempt to come up with the right answers is mostly a distraction . because generally speaking, there is no %ight+answer.

This is no secret within the retirement planning community, but perhaps we do not take it seriously enough. For if the voluminous academic work over the last decade or two on retirement income planning has proved anything, it's that even when a person is going through the retirement process, the adequacy or inadequacy of their savings is something of a crapshoot, which is why Monte Carlo analysis is an intuitively appealing approach to that problem. So take that reality, and multiply it by the

unknowns that someone at, say, age 40 is still facing *before* retirement. Between that age and retirement, most working people will change jobs at least once, most will move at least once, many will have additional children, will divorce or marry (or both, sometimes more than once), will develop chronic illnesses or suffer disabling accidents (or a spouse will), and most will have other unexpected changes in their pre-retirement lives that will have major effects on their finances and/or their life expectancy.

Given all that, the idea that any calculation, simple or complex, can reliably come up with the *right*+retirement savings target has no merit at all in most situations, and even in the best of cases is subject to significant uncertainty.

Furthermore, *even in theory* there is no *right*+number. Even after the fact . when the consumer and all his or her family has died and all the financial outcomes are determinable . there is no solid basis for saying what the right amount of savings at retirement age would have been. This is partly because financial needs are highly elastic . you can say once they're gone what someone *did* spend, but no one can say how much they *should* have spent to optimize their welfare and their happiness. But it is also because, even if you could identify what turned out to be the optimal retirement savings after the fact, it generally would not have been prudent for that person to save exactly that amount, because the outcome *might* have been different. Retirement savings have to cover contingencies, not just realities.

So seeking to come up with the *right*+retirement savings number is chasing an illusion, and therefore is not a proper objective for rational analysis. Unfortunately, the reason consumers want to use these tools is because they think that *the* number exists, and no one really wants to discourage them from going through the process. But that is a marketing issue, and it can be worked around.

When designing planning methods and tools, knowing that it's impossible to accomplish the ostensible objective of these calculations frees us to focus on their actual objective: to motivate savings.

But that isn't what has happened so far. Most retirement gap calculators were built with the impossible objective foremost in mind, so they are failing both at that objective (which is unavoidable) and at the true objective (which is very much avoidable).

The current methods work decently as a motivator only for people whose need for additional savings is not very great. Although this can be true for people at any age, depending on their circumstances, in general it's most commonly the case for people still three or four decades away from retirement. It can, in fact, be a powerful incentive for those individuals to save. Start at age 25, sock away perhaps \$240 a month, and kazaam! your retirement is provided for.

But the closer you get to retirement, the more discouraging it is to try to close the gap. For example, to retire at age 65, replacing 70% of an \$80,000 annual income for 25 years at retirement, assuming \$20,000 in Social Security and zero current savings, requires \$460,202 in new saving. Assuming a 6% annual return and no inflation adjustments, monthly savings needed at different ages equal:

Monthly savings required	
Age 25	\$ 240 / mo.
Age 30	\$ 333 / mo.
Age 35	\$ 470 / mo.
Age 40	\$ 677 / mo.
Age 45	\$1,010 / mo.
Age 50	\$1,596 / mo.
Age 55	\$2,819 / mo.
Age 60	\$6,591 / mo.

In this reasonably typical case, waiting five years (until age 30) increases the required level of savings nearly 40%, and if you wait until age 60, you need to save, in this case, nearly 100% of your pre-tax income!

One result is a potentially powerful argument for 20-somethings to start early . but one that often falls on deaf ears, either because retirement seems too far away to care about or because young people have other legitimate financial priorities.

Another and very unfortunate result is a powerful *disincentive* for people close to retirement, or even ten or twenty years away. Of course, these numbers are smaller for some people, especially those who have good pension plans at work and/or who have in fact started saving earlier. But even for those people, getting close to retirement means that the required savings shoot up astronomically, simply because there is not enough time left to fill whatever gap may remain.

The normal reaction to such numbers is: why bother? People decide they will either live on whatever they have, or they will hope for a miracle: winning the lottery, or maybe even dying before their money is gone. Or they just decide they will never retire at all . not realizing that only a minority of people actually have this option (at some point, they will lose their job, or they will have to leave because of illness or incapacity or perhaps to become a caregiver for a spouse, or for some other reason).

Unfortunately, we have no solid data linking all these factors. But studies pretty consistently show that, especially among Baby Boomers, somewhere around half of the population is ~~at~~ **at risk** of having an inadequately funded retirement,¹ and, probably not coincidentally, about half of the population is not confident about their retirement security.² Frighteningly, a recent AARP survey found that close to half (44%) of respondents said that they think they may never be able to afford to stop working.³ It is almost certainly the case that many people have fallen to this level of despair because they have not been shown in an effective manner how an acceptable level of current preparation can make an important difference in their retirement readiness .

¹ For one of the most recent analyses, see Jack VanDerhei, *What Causes EBRI Retirement Readiness Ratings™ to Vary: Results from the 2014 Retirement Security Projection Model®*, February 2014, available online at: http://www.ebri.org/pdf/briefspdf/EBRI_IB_396_Feb14.RRRs2.pdf

² For example, *Redefining Retirement: The New 'Retirement Readiness' - The 13th Annual Transamerica Retirement Survey*, May 2012, available online at: <http://www.transamericacenter.org/retirement-research/retirement-survey>

³ AARP *Retirement Attitudes Segmentation Survey*, 2013, available online at: http://www.aarp.org/content/dam/aarp/research/surveys_statistics/econ/2013/Retirement-Attitudes-Segmentation-Survey-AARP-rsa-econ.pdf

but on the contrary, they have been discouraged by simplistic assumptions, either of their own or of the poorly designed tools they have used to calculate their needs.

What can be done to rectify this situation?

Solution #1: Fudge the assumptions

Here's one benefit to the unpredictability of the future: when you're doing projections, you are at least somewhat justified in *inventing* and illustrating the future you'd like to have. Or to put it more plainly, you can fudge your assumptions, adjusting them until you get the answer you want, or at least an answer you can live with.

It's easy and fun, if you have the right kind of software. Back in our earlier example, which could have been made even more dire, we actually already made *three* favorable assumptions: (1) that savings would be amortized over the expected lifetime, while in real life most retirees prefer just to live off the interest; (2) that income is not increasing *before* retirement, because inflating pre-retirement income moves the goalposts further away; and (3) that living expenses do not inflate *after* retirement. Now let's see what happens when we make additional favorable assumptions:

<u>Assumptions:</u>	Monthly savings required	
	<u>40-yr-old</u>	<u>60-yr-old</u>
Same as before	\$677	\$6,591
Delay retirement by 5 years	\$470	\$2,819
8% return on investment (was 6%)	\$420	\$5,235
20-year lifespan in retirement (was 25)	\$608	\$5,914
Income replacement at 60% (was 70%)	\$527	\$5,126
Increase savings 5% a year	\$405	\$5,998
<u>All five of the above changes</u>	\$225	\$1,243

Clearly, the numbers we get by changing the assumptions are better (i.e., less frightening). Most are not vastly better, though. Also note that the impact of the assumptions varies depending on the client's age. For the younger client, the biggest impact comes from using an *increasing rather than a level savings* pattern, followed closely by *increasing the assumed rate of return* by 200 basis points. For the older client, *postponing retirement* by five years has by far the biggest impact, roughly equal to the other four combined. (The magnitude of these effects, of course, depends on all the variables, not just these, so we shouldn't read too much into this one example.)

My own experience, which embraces over 25 years of working with financial calculators of this general kind, is that consumers are fairly comfortable with fudging the assumptions. They would much rather do that (and embrace hope) than face reality (and actually sacrifice in some way). When there's a financial advisor of some kind, however, fudging is more problematic. The risks can be explained during the process, but fudging can make the advisor look bad (dishonest or callous) and also is likely to produce results that will later boomerang on the client and on the advisor.

So Solution #1 has to be used judiciously, but it should be an option. Unfortunately, most existing gap analysis tools are too rigid to permit this solution in the most palatable ways. In an effort to remain uncomplicated, they fail to provide sufficient options to adjust the assumptions. If just one or two or three assumptions can be modified, only a minority of clients will be helped. Different consumers are comfortable with

different options, and different options have different effects in different circumstances. So tools that provide more ways to manipulate the calculations offer the highest probability of producing a motivating outcome. Failure to recognize this reality has made the financial industry far less effective in encouraging retirement savings than it could be, which ultimately is a disservice to their customers as well as to themselves.

Even so, added computational flexibility is not enough. For people not already in the ballpark of being able to afford the retirement funding they need, we also require:

Solution #2: Change the thinking behind the model

The main reason that gap analysis is least effective for the middle-aged and older consumers who both need it the most and are in the best position to save, is that they are already close to the edge of the cliff. They have fewer years to add to their nest-egg, and less time for the miracle of compound interest to help them.

But what if we could change the game, so that the cliff was farther away? What if instead of the cliff occurring at retirement, it could be postponed until death? And what if this were not some kind of trickery, but it was actually the most sensible and justifiable way to think about retirement?

All these what ifs are true . . . and there's a good reason why. It has to do with the idea of consumption smoothing that goes back to the 1950s and a paper by Modigliani and Brumberg and a book by Milton Friedman⁴, later applied by Laurence Kotlikoff to retirement planning modeling.⁵ In simple terms:

***What matters most is not what you save but what you spend.
Mathematically, these are the same. What you save equals
whatever you don't spend.***

***But in gap analysis, saving and spending look different for an
artificial reason: you can save only while you are working, but
you can spend less for your whole life.***

It may seem contradictory to say that saving and not-spending are the same, while also saying that shifting the focus from the first to the second would make a mathematical difference. But it's not, because there is a fallacy in traditional gap analysis, namely, that the new savings the client takes on have no effect on the percentage of income the client needs to replace in retirement.

If this seems a bit confusing, let's look at a relatively extreme example: our 55-year-old from the first table of numbers needs to save \$2,819 a month, or nearly 40% of income. In order to do this, the client needs a simpler lifestyle and lower expenses. But if these cuts occur, and the client can live on this, why assume that they disappear at retirement? If the client can live on less, and gets used to it, doesn't it make more sense to assume that these adjustments continue indefinitely?

⁴ Modigliani, F. & Brumberg, R., 'Utility Analysis and the Consumption Function: An interpretation of Cross-section Data'. In: K.K. Kurihara (ed.): *Post-Keynesian Economics*, 1954. And Milton Friedman, *A Theory of the Consumption Function*. Princeton N. J.: Princeton University Press, 1956

⁵ Laurence J. Kotlikoff, 'Economics' Approach to Financial Planning,' *Journal of Financial Planning*, March 2008, available at <http://www.esplanner.com/learn/economics-approach-financial-planning>.

If we focus on this change in lifestyle that increased saving represents, and assume that that lifestyle persists rather than pops back up in retirement, *we can stretch the period of “saving” through the entire lifespan, not just until retirement age. This can greatly reduce the size of the hit to lifestyle in the first place, and thereby reduce the scariness of saving for retirement.*

So if we go back to our original test case assumptions, we end up spreading the cost of the ~~%savings+~~ ~~%lifestyle adjustment+~~ over 25 years more than we do using traditional gap analysis. As you can see, the impact is significant across the board, but it is increasingly dramatic as people get closer to retirement.

<u>Assumptions:</u>	Monthly savings required	
	<u>Traditional Method</u>	<u>Solution #2</u>
Age 25	\$ 240 / mo.	\$ 222 / mo.
Age 30	\$ 333 / mo.	\$ 299 / mo.
Age 35	\$ 470 / mo.	\$ 405 / mo.
Age 40	\$ 677 / mo.	\$ 549 / mo.
Age 45	\$1,010 / mo.	\$ 750 / mo.
Age 50	\$1,596 / mo.	\$1,030 / mo.
Age 55	\$2,819 / mo.	\$1,431 / mo.
Age 60	\$6,591 / mo.	\$2,017 / mo.

In case you're wondering about the impact of using *both* solutions . fudging the numbers *and* using a consumption-smoothing approach:

Combining Solutions #1 and #2 reduces the monthly saving requirement for our Age 60 client to only \$653, which is less than 10% of the initial determination of \$6,591.

Having the ability to perform retirement gap analysis in both of these ways is a complete game-changer . especially for Baby Boomers who are still working but are rapidly approaching their desired retirement age without having saved what they should have.



Let's keep in mind that the examples presented here are illustrative . individual circumstances vary considerably. Even so, the ability to analyze closing the ~~%gap+~~ as a lifelong rather than a temporary adjustment to the household budget will have a favorable impact on the acceptability and motivational force of a retirement analysis in almost any circumstances. In general:

For people in their 20s and 30s, savings requirements usually illustrate as fairly low numbers anyway. However, incomes are also relatively low, and especially for young families, budgets can be tight. The ability to illustrate a lower monthly commitment is often essential to inspiring such people to develop the habit of regular savings.

As they move through their 40s and 50s and often beyond, it becomes harder and harder to close the gap. Fudging the assumptions enough to make the medicine palatable becomes increasingly difficult, and eventually impossible. But extending the savings into retirement, and not just until retirement, can persuade these clients that there is still the time and the means to rectify their retirement outlook.

Implications for retirement gap analysis calculators

Taking all of this into account, the truly beneficial retirement gap calculator is clearly different from most models in use today. It is still simple to use, but it also allows the user to tweak the assumptions in a significant number of ways, including the option of using a consumption smoothing algorithm to spread the impact of %savings+ through the expected lifetime, and not just during the expected working years. As our examples show, these capabilities enable more-or-less realistic savings needs to be illustrated across a wide spectrum of actual dollar commitments. Using such a calculation tool, most consumers (and any financial advisor) could come up with a result that is dramatically less de-motivating even in cases where the initial numbers are overwhelmingly intimidating.

Currently, such calculators are hard to come by. In particular, the models in the marketplace that best illustrate consumption smoothing (Kotlikoff's ESPlanner⁶ and my own company's RetirementWorks II⁷) are too sophisticated to qualify as %simple+ tools. Both have enjoyed only modest commercial success over a period of a half dozen or more years, despite being analytically far superior to other models, and the reason is almost certainly their failure to meet the financial industry's desire for easy-to-use tools.⁸

Fortunately, if you are looking solely at retirement savings, you can get the effect of consumption smoothing, as previously noted, simply by extending the savings time period until death, rather than just until the retirement date (though this option would

⁶ <http://www.esplanner.com/>

⁷ <http://www.retirementworks2.com/>

⁸ I cannot speak for ESPlanner, but with RetirementWorks II we have found that it's the financial professionals who have no patience for more sophisticated models. Consumers ó at least those who are close to retirement or already there ó clearly prefer deeper analysis. Our paying customers can choose between a moderately detailed process and a very detailed process, both options at the same financial cost, and almost without exception they choose the more detailed option. Financial professionals, by contrast, almost always choose the least detailed option available.

have to be built into the software, which almost never is the case). One caveat, though: people who are already taking a significant amount out of their pay to save for retirement may already be assuming a lower income replacement ratio to reflect this. Adding a simplified consumption smoothing algorithm on top of this is double-dipping, and therefore unfairly low-balling the savings required. So when consumption smoothing is being used, the replacement ratio should not also be held very low, unless there are separate reasons for doing so. ***Specifically, the replacement ratio should be the ratio between anticipated post-retirement and pre-retirement spending, not income.***

Conversely, of course, you can simulate the effect of consumption smoothing in a traditional, simple calculator just by reducing the income replacement ratio. But doing it that way requires that someone manually calculate, or (worse) estimate, the consumption smoothing effect, so this is at best an awkward alternative.

In this context, it is helpful to emphasize that there is an educational component to the motivational objective of retirement gap analysis. The idea is not simply to pump people up emotionally about the need to save, but rather to demonstrate to them what they ought to do, and to some extent, why they ought to do it that way. In the end, it's about persuasion, not enthusiasm.

That's why the ideal model allows the software user to modify a fairly large assortment of assumptions. A knowledgeable financial advisor can walk them through the process, and show them the bottom-line effects of different assumptions, and can discuss which assumptions can more safely be adjusted (based, for example, on the consumer's degree of control over those factors). A consumer using such software on his or her own can achieve some of the same effects by trial and error. Well-designed software can also give users guidance on what kinds of changes will have what kind of effect.

Being able to see these effects educates the consumer about which changes are mostly likely to make a difference in the most reasonable and most palatable ways. For example, is it better to assume a higher level of saving, or a delayed retirement date? Up to some common-sense limit, the more such options are available to the software user, the more readily the software can produce a result that has sufficient computational plausibility and sufficient doability to motivate the consumer to actually make the commitment to less spending and more saving now.

This takes pencil and paper worksheets out of the running. Software is required, but it needs to be made more sophisticated (in a limited and prudent way) than almost all such tools are currently designed to be. Otherwise, we will continue to promote among a large swath of consumers a sense of despair about retirement instead of demonstrating a realistic path to making a decent retirement possible.

Charles S. Yanikoski, President

Still River Retirement Planning Software, Inc. / RetirementWORKS, Inc.
www.StillRiverRetire.com / www.RetirementWorks2.com (978-456-7971)

One working example of a calculator of the kind we describe can be found at:

<http://www.stillriverretire.com/newcalcs/RetNeeds.asp>.

However, this is a demo version, not intended for regular commercial use.