NYU Wagner

Robert F. Wagner Graduate School of Public Service

Adapting Social Security Policy for the Long Term

September, 2007

Steven W. Popper[1]
Senior Economist, RAND Corporation

Introduction

From the perspective of policy analysis, three groups of factors affect the ability of Congress to legislate wisely for the long term. The first are those related to Congress's institutions, processes and political interests that have been detailed in other contributions to this series. The second is the difficulty analysts face in providing definitive findings to policymakers as the time horizon lengthens. The third is the difficulty of meshing the styles and requirements for rigorous analysis with the needs and language of legislators and decision makers. The future of Social Security[2] provides the perfect storm conditions that bring all these into view.

Although this brief is based on current research exploring how to resolve the difficulties raised by this intersection, it is not a report of that research.[3] No prescriptions on specific policies will be given nor any claims made that the analysis itself is definitive. Rather, as a proof-of-principle illustration (and of only one aspect of the larger Social Security issue at that) it is intended to make concrete and demonstrate the viability of a shared vision of process for assisting Congress that is implicit in several briefs that have previously appeared in this series. This vision is worthy of serious consideration in any program to enhance Congressional capability to legislate for the future.

The discussion begins by specifying the problem faced by policy analysis when confronting deep uncertainty[4] and why this poses serious obstacles to any deliberative body such as the U.S. Congress. The brief then outlines an analytical method designed to be practical while making explicit that which everyone knows: we cannot be sure what the future will bring. Instead of deciding a priori what the "best guess" future might be and then choosing policies to optimize based on those conditions, the method systematizes the type of inductive thinking that responsible and dedicated legislators and policymakers actually engage in by asking

The Legislating for the Future Project is an initiative of New York University's John Brademas Center for the Study of Congress and the Organizational Performance Initiative, and is co-sponsored by the Brookings Institution and the RAND Corporation. The project will examine the capacity of Congress to address long-term problems facing the nation, probe the public's attitudes towards Congress' ability to make long-term decisions for the 21st Century, and analyze specific long-term policy issues. The Legislating for the Future Project will convene experts for discussions of specific long-term issues, such as global warming, and seek to generate strategies to make Congress more flexible and adaptive to future problems. The Advisory Committee for the project is headed by Former Representative Lee H. Hamilton. The project is funded by the John Brademas Center for the Study of Congress, the Smith Richardson Foundation and the Carnegie Corporation. For more information, please visit: www.nyu.edu/wagner/performance and www.nyu.edu/brademas.

tacitly or out loud, "What if..?" The discussion then describes briefly a reasoning process that is iterative, interactive, and inclusive of different classes of knowledge, expertise and perspectives on the world. The object is to reveal potential solutions that perhaps may not be optimal for any particular set of assumptions or values but will nevertheless meet minimal criteria for acceptability, set by the political process, across a wide range of the plausible futures we may well confront. The brief concludes with thoughts on how incorporation of robustness thinking and explicit recognition of uncertainty may affect legislative processes.

While the focus of the following discussion is Social Security, the true subject is how to conduct an analysis of policies that play out over the long term. Exchanging our usual question "what is likely to happen?" for a better question ("given that we cannot reliably predict what will happen, what is our best course for the short term?") is a subtle transformation that permits systematic, quantitative analysis of an issue notoriously resistant to such treatment. Perhaps most important, it creates an avenue for framing policies that are designed from the outset to be both flexible and adaptive, a crucial value when legislating for the future.

ANALYSIS OVER THE LONG TERM OF LONG-TERM POLICY

The title of this brief would seem redundant. After all, Social Security, by its very nature is an instrument designed to operate over the long term. Current workers' contributions are not banked in a conventional sense. Rather, they are used to pay the benefits of current Social Security recipients. Later, those who contributed earlier are paid back from those who are then currently in the work force. By design, the system is inherently intergenerational and long term.

Yet, when we consider measures to ensure the ability of this system to match its revenues and savings to current and future entitlements ("solvency") we face a paradox. The system extends over the long term yet our ability to perceive and act over great expanses of time are quite limited. Social Security solvency is but one example of a type of issue that consistently confounds our ability to analyze and frame wise policy. For illustration we need look no further than the 2006 (or, indeed, any) annual Social Security Trustees' Report.[5]

Though only required by law to look out five years, the Trustees wisely choose to look out over a considerably longer term to scan for possible shortfalls. They possess data on demographics, statistics on insurance utilization, and sophisticated models and forecasting tools. But they do not possess perfect foresight. Small shifts in assumptions well within the range of plausibility produce dramatically different conclusions over decades. Using variable values to represent what the Trustees term the Intermediate Cost case, the trust funds[6] set aside to cover Social Security

obligations become exhausted in the year 2040. At that point, any resources to cover additional obligations would presumably have to come from general revenue sources if benefits were not to be adjusted to fit within the means available. If the High Cost assumptions are applied, the exhaustion would occur a decade earlier, in 2030. On the other hand, utilizing the Low Cost assumptions, the trust funds never go negative and there would be a positive and growing balance of nearly $8 trillion in the year 2050.

Clearly, these three forecasts also have three different implications for policy. One suggests staying the present course while the others suggest the need for short-term attention while small changes may still suffice to correct our course. What makes this even more problematic is that the range of outcomes spanned by these High, Intermediate and Low Cost cases actually cover only a relatively narrow band of the range of plausible futures. Plausible changes to certain key variables would have the power to place the future that actually transpires well beyond either the High or Low Cost cases framed by the Trustees Report. The Trustees and their analysts would be among the first to acknowledge this fact.

In this sense, our usual means for analysis appear inadequate to help Congress reliably evaluate what courses of policy to consider. Generating multiple, credible forecasts is easy and the Trustees' Report is at great pains to do so in as sophisticated a manner as present art allows. What they cannot provide are reliable predictions. Legislative deliberation is stymied if political discourse becomes an irresolvable debate over which presently unknowable set of future values are the correct assumptions for the purpose of planning short-term policy. Worse yet may be a state where the Trustees' Intermediate Cost case, based on one set of assumptions, is taken as "truth" and the debate is over how to handle the resulting insolvency. The addition of ideological and political concerns completes a recipe for friction, inaction or ill-advised policy actions.

ROBUST DECISION METHODS AID POLICY CHOICE

A recent general audience survey article on economic modeling raised troubling questions for model builders and the consumers of model output.[7] It portrayed a Red Queen's race in which ever more sophisticated generations of models seek enhanced predictive capabilities – only to see the prize consistently remain out of reach. Yet, even if such models were considerably more complicated than they currently are they still could not generate reliable predictions for even just a decade from now. That being the case, it is worth asking what use any

model can be for illuminating policy choices over issues as complicated as those embraced in the Social Security solvency problem.

The root of the apparent problem lies not in the models themselves but in how they are used. Most models are created to play a role in a "predict-then-act" process of analysis: The steps are to first develop the most accurate possible model of the system of interest, gather data and make assumptions, generate predicted outcomes, and then apply the tools of optimization to find the best course for action. Is this a reasonable use of a model, however, when prediction is not credible? If we have optimized for one, supposedly most-likely, future and the actual future turns out differently, we can hope that our previously optimal plan will still be serviceable but we have no proof that this will, indeed, be the case. We might find ourselves bound on a course that is clearly deleterious to our interests given the way the state of the world has changed from what had been expected.

Yet, predictions themselves are rarely what we seek. Rather, what we would wish to understand is how changes in the future might affect our choice among alternative actions today and how the actions we do take will affect our chances of being successful in meeting our goals. Our true interest, once we acknowledge that we cannot be sufficiently predictive, is to understand how we can choose today's actions most wisely in light of our long-term objectives.

The shift in focus from model forecasts to informing decisions is subtle yet resolves many conundrums. Instead of determining the "best" model and solving for the strategy that is optimal (but fragilely dependent upon assumptions) we should instead seek among our choices those actions that are most robust – that achieve an agreed level of goodness across the multiple models and assumptions consistent with known facts. This is closer to the actual policy reasoning process.

What we need from a model is not a prediction. Rather, a model serves as an artifact that contains what we understand about critical relationships among key factors and that can then be used to generate the myriad scenarios of the future that are consistent with our current information. As we systematically vary assumptions about factors whose future values are presently unknowable, we generate an ensemble of alternative futures -- a test bed for helping select among policy alternatives. Rather than characterizing uncertainties at the beginning of the analysis either by assigning values, assuming probability distributions, or dropping them entirely pending later analysis, we leave the uncertainties uncharacterized in terms of probabilities but nevertheless explicitly represent them in the model. The focus of the analysis then becomes not what assumptions we should choose but rather what conditions we would need to believe were likely in order to favor policy

"A" over policy "B" – and how we might construct a policy "C" that may relieve us of the need to choose.[8]

HOW CAN ANALYSTS AID LONG-TERM DECISIONS?

It is convenient to map factors we deem important in the analysis into four categories:

> ***X*** (or e***X***ogenous) factors outside our control that may affect outcomes in the future and render some strategies superior to others in retrospect after a decision has already been taken;
>
> ***L***evers constituting actions under our control that may either be combined with others by variations in composition, degree and sequencing into alternative strategies or else explored individually on their own;
>
> Cause-and-effect ***R***elationships between the actual state of the world (represented by those characteristics explicitly explored as ***X*** category factors) and actions we may take (***L***) that yield the outcomes we wish to measure (***M***); and
>
> ***M***easures for assessing whether outcomes resulting from taking actions (***L***) within a particular environment (***X***) meet our criteria for goodness or not. This is a category for exploration because no single measure may be sufficient to satisfy the criteria for goodness held by the parties to a decision. Similarly not everyone will agree on what weights+ to place on different measures of outcome success or failure.

The XLRM framework serves several purposes. It is a transparent check list that will allow others to observe with precision what unknowns and assumptions are being modeled. It provides an intellectual bookkeeping system permitting further refinement by parties to the analysis of their own perception of the problem and issues to be addressed. Perhaps most of all, it serves as a template for design of the analytical tooling that will be used to support the investigation of alternative policies under a variety of emergent conditions.

Table 1 presents the XLRM design being used in the RAND study. It is quite limited in scope and so offered only as an illustration. It is framed around solvency, by no means the only important Social Security issue. A full policy analysis of Social Security, even if only limited to the solvency issue, must contain a richer set of elements in each quadrant before its findings could have credible implications for policy.

Possible Future States of the World ("***X***")

The academic fields of demography and economics both wield powerful tools that often succeed in identifying trends and predicting

outcomes. However, the longer the time period, the less reliable are the findings from these two disciplines.

The 2006 OASDI Trustees' Report recognizes the difficulties of being accurately predictive. Even in the medium term, looking out to the year 2020, differences in estimates appear. The Intermediate Cost case results in estimated balances in that year of $3.64 trillion while the Low Cost and High Cost cases yield estimates of $4.91 and $2.22 trillion, respectively. Therefore, just fourteen years out the Trustees' estimates vary by more than 120%.

Uncertainties Outside Control (X)	***Levers Under Control (L)***
• Average children per woman • Annual net immigration • Mortality changes • Labor force participation, M/F • Disability incidence and recovery • Unemployment • Inflation • Productivity growth • Wage share growth • Annual average hours worked • Real interest rate	• COLA rate modifications • General fund revenue transfers • Date of strategic reassessment • Policy change rules
Relationships Between Factors (R)	***Measures of Outcomes (M)***
• macro model of aggregate program finances • micro model of cohort individuals (both included in the SSASIM© software package)	• OASDI balance in 2020 • OASDI balance in 2050

Table 1. Selected Factors for Exploring Social Security Solvency Actions and Goals

This widening of plausible futures becomes greater the farther out one looks. It is one reason for the tendency to restrict possible variations to a somewhat narrow range.[9] However, if we wish to reason across ensembles of future states of the world, those worlds must be sufficiently inclusive to provide us with reasonable assurance that the actual future that will transpire is included within the ensemble. We wish to make use of the logical principle that whatever true statements we can make for an entire set must therefore be true for any particular member of that set. If both our landscape of plausible futures is wide and we can find actions

that appear to be robust broadly across that landscape, then those actions are likely to be well-suited to whichever of those plausible futures turns out to be the one we face.

It is also important to make certain all relevant assumptions about different futures are represented. This way groups and legislators on opposite ends of a debate will recognize that their views are being accorded the same validity as any others.

		RAND			Trustees (2006)	
DEMOGRAPHIC UNCERTAINTIES	**Min**	**Nominal**	**Max**	**Min**	**Nominal**	**Max**
Ave. children per woman	1.4	2.0	2.8	1.7	2.0	2.3
Net immigration (millions/year)	-0.5	0.9	4.0	0.7	0.9	1.3
Mortality rate decline (%/year)	-1.0	**0.0**	1.0	0.3	**0.7**	1.2
Female labor force participation	40%	61%	75%			
Male labor force participation	40%	74%	80%			
Disability incidence factor	0.9	1.1	2.0			
Disability recovery factor	1.0	1.3	2.0			
ECONOMIC UNCERTAINTIES						
Unemployment rate (%/year)	3.5	5.6	10.0			
Inflation rate, based on CPI (%/year)	-0.3	2.8	6.0	1.8	2.8	3.8
Productivity growth rate (%/year)	0.0	**1.6**	3.0	1.4	**1.7**	2.0
Wage share growth rate (%/year)	-1.0	-0.2	1.0			
Hours worked, ave. growth (%/year)	-1.0	0.0	1.0			
Real interest rate (%/year)	0.0	**2.0**	4.0	2.1	**2.9**	3.6

Table 2. Range of Values Used in RAND Analysis and by OASDI Trustees

Table 2 illustrates ranges of future values explored across thirteen presently unknown variables. The three left-hand columns of figures indicate the minimum and maximum values used in RAND's analysis as well as the nominal values that initialize the analytical system.[10] For purposes of comparison, the three right-hand columns show the implicit ranges utilized in the Trustees' Report, rephrased in this instance to

represent minima and maxima. The "nominal" case values ascribed by us to the Trustees' Report are actually those values used by them to construct their Intermediate Cost case. The ranges used by RAND are wider than those implicitly employed by the Trustees in most cases.

Policy *L*evers: Alternative Actions to Maintain Solvency

The RAND study is intended as a demonstration and not a conclusive study of policy conclusions. Four competing strategies were chosen to reflect important aspects of the current debate without corresponding to any specific proposals. In each scenario simulation run, one of these actions is selected in 2006 and employed until 2020. In that later year, the current state of the OASDI trust funds is assessed. If the result is deemed unsatisfactory for a particular scenario run, the current strategy is abandoned in favor of one of the other three depending upon the severity of the perceived or anticipated crisis in funding.[11] In this way, each strategic course includes the simplest possible adaptive strategy where an indicator triggers differential action at a later date depending upon its actual value at the time. Clearly, much more sophisticated adaptive strategies are possible and would be explored in a fuller policy analysis.

All strategies besides **"Hold Steady"** introduce changes into Social Security policy. **"**Hold Steady**"**, however, maintains policies current in 2006 through the year 2020. Depending on the state of the OASDI trust funds' balance in that year, the "Hold Steady" policy for that scenario simulation run either continues or, depending upon how severe the decline in trust funds' levels proves to be, there is a shift to either the relatively mild "General Revenue Transfer" or more aggressive "Adjust+Transfer" policies for the remaining years until 2050.

"Cost-of-Living Adjustment" enhances solvency by reducing benefits. In contrast to current policy that mandates an annual COLA to Social Security benefits based upon the changes in the consumer price index, the strategy reduces that annual adjustment by a 0.5% decrement each year.[12] The status of the OASDI trust funds in 2020 would mandate either continuing the COLA policy, reverting to the "Hold Steady" course, or applying the more stringent measures included in the "Adjust+Transfer" policy.

While the COLA policy affects the benefits side of the solvency equation, the **"General Revenue Transfer"** (GRT) strategy operates on the revenue side. Revenue would be transferred from the federal general revenue accounts to the OASDI trust fund accounts at a fixed annual rate.[13] In this way, the funds could benefit from the interest earned during the years of surplus to delay or avoid possible shortfalls in later years when the costs of the Social Security programs exceed the receipts from OASDI

taxes. We chose the rate of 0.5%.[14] In 2020 a decision is made to revert to "Hold Steady", maintain the GRT policy, or employ the "Adjust+Transfer" policy.

"Adjust+Transfer" addresses the revenue and benefit sides simultaneously. In this case, both are pursued with greater vigor than in either the pure revenue (GRT) or benefit (COLA) cases. The arithmetic COLA decrement becomes a full 1.0% taken off the top of the CPI and the revenue transfer rate is raised to 1.0% as well. As before, 2020 is the year that policy shifting may occur.

Clearly, only a narrow range of actions are explored in this study. Others might include means testing and scaling Social Security benefits; diverting a portion of OASDI tax receipts into private accounts; changing the retirement age, payroll tax rate or ceiling; and reducing disincentives to continuing to work, among others. Further, hybrid strategies could be constructed from several actions and by varying rates and thresholds. {This concept will be discussed further below.}

***R*elationships: SSASIM©**

As both a practical matter and a rhetorical strategy we used an existing simulation model obtained from Martin R. Holmer of the Policy Simulation Group, SSASIM©[15]. SSASIM is currently the sole scenario generator in the "***R***" quadrant in the XLRM matrix. This means the explorations in the analysis are of parametric uncertainty – changes in values of presently unknown input variables, policies or measures of merit. This need not be so. To the extent to which there are competing models or alternative conceptions of causality, these implicit structural uncertainties may also be made the subject of an analysis of the type we present.

Measures

In our current study we look only at OASDI trust fund balances in 2020 and 2050. This is a very narrow perspective on the Social Security debate in the U.S., deliberately so given the illustrative purpose of this work. Other measures would be urged by different partisans to the solvency debate as being as, or perhaps even more, important. This highlights the crucial point that the choice of measure is as much a political as an analytical decision, if not more so. Candidate robust strategies should be tested for the ability to satisfy all these criteria as well.

A METHOD FOR FRAMING AND TESTING POLICIES AGAINST THE UNKNOWN

To discover courses of action that are demonstratively well-hedged against uncertainty while satisfying our criteria for acceptability, we first generate a landscape of alternative states of the world. This is crafted

from the thirteen uncertainties we identified in the "X". The analysis uses a sample of 4,000 futures as the test set. [16]

Observe Policy Performance Across Scenarios.

We may now compare the performance of the four strategies. Simulations are generated by running each policy lever in each of our 4,000 alternative futures and assessing outcomes.[17] This yields an ensemble of 16,000 different scenarios.

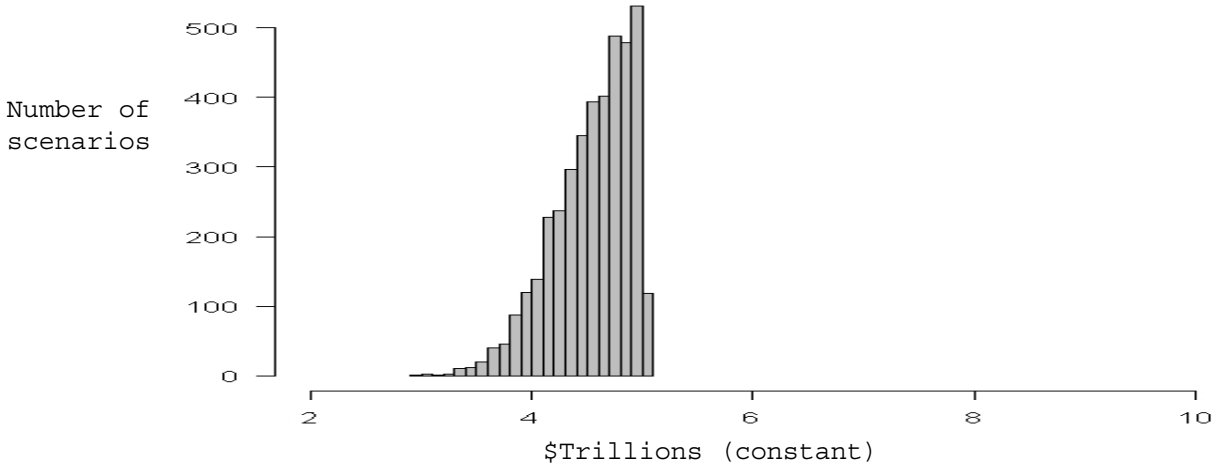

Figure 1. OASDI Balances in 2020 for "Hold Steady" Policy

We then derive from these results systematic insights into the relative strengths and weaknesses of each strategy. Figure 1 is an aggregate view showing the resulting balances of the OASDI trust funds in the year 2020 if we apply the "Hold Steady" policy in each of the 4,000 alternative futures. The horizontal axis shows the trust fund balances in trillions of constant 2006 U.S. dollars. The vertical axis shows the number of scenarios that fall into each $100 billion-wide band. For comparison, the same view for the "Adjust+Transfer" policy is shown in Figure 2.

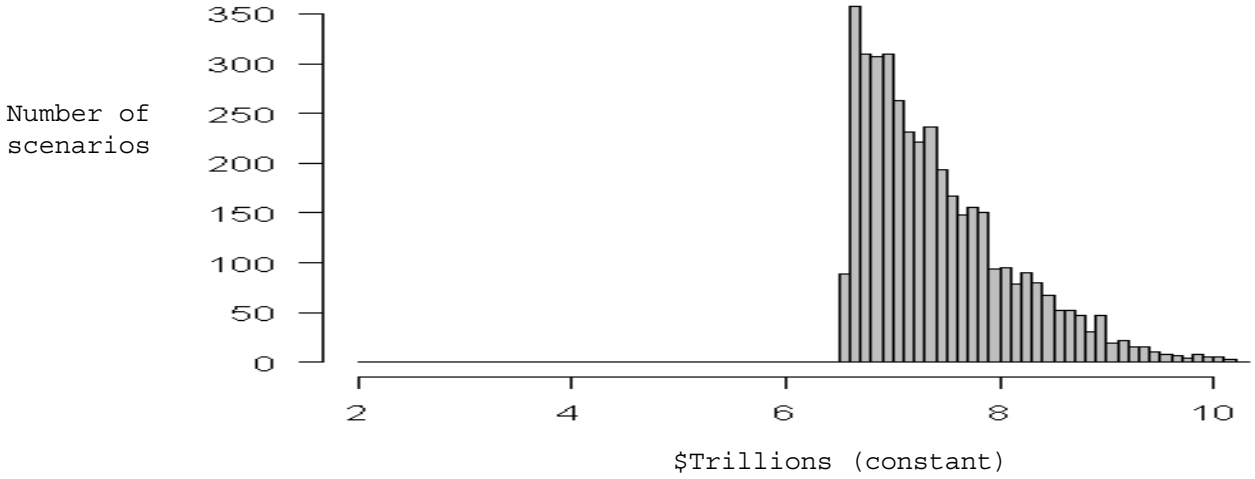

Figure 2. OASDI Balances in 2020 for "Adjust+Transfer" Policy

Not too surprisingly when looking out only 14 years, solvency is not an issue. Looking across all 16,000 scenarios, the state of the OASDI trust funds balance in 2020 may be as low as $2.9 trillion for some harsh futures where we continue to "Hold Steady" to over $10 trillion when conditions are more favorable but we have applied the "Adjust+Transfer" policy is applied.

Looking out to 2050, Figure 3 shows more dispersion for all 16,000 scenarios.[18] Sixty percent of the scenarios for all four policies report trust funds' deficits with the mean at -$770 billion. Clearly, many scenarios show results quite wide of the mean in both directions. Some scenarios show positive trust funds balances of over $50 trillion. This wide variation is a function both of the power of small differences to bring about large changes over time and the relatively simple-minded nature of the four strategies policies.

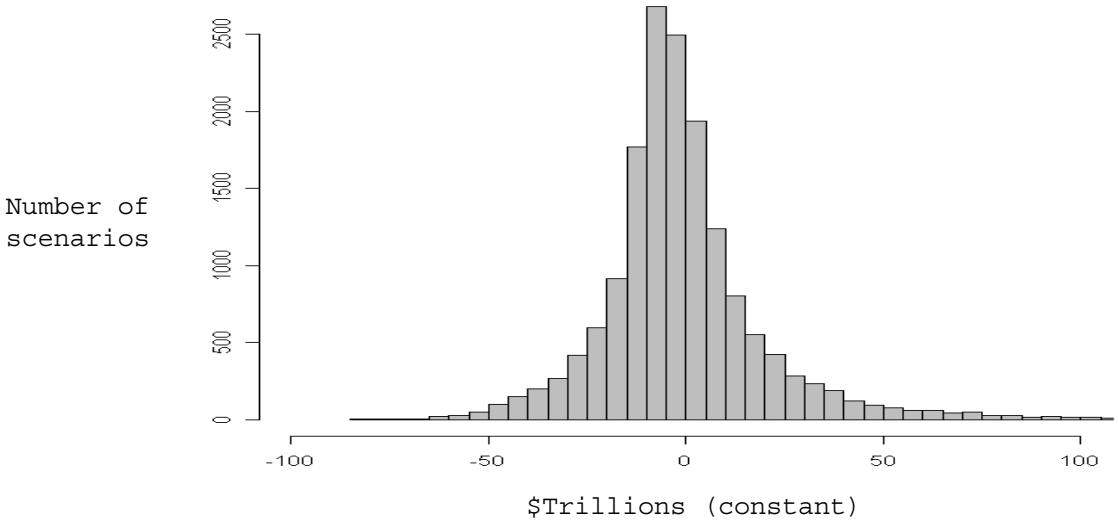

Figure 3. OASDI Balances in 2050 Across All 16,000 Scenarios

By comparison, the Trustees' Report projection to 2050 is $7.92 trillion for their Low Cost case.[19] The aggregate projections of the 16,000 scenario runs are somewhat more pessimistic than those used by the Trustees, if we may assume that their Intermediate Cost case point estimate may be taken as a mean value. This is not surprising; SSASIM contains a less sophisticated model of economic response to changing conditions than the models used by the Trustees. There are fewer endogenous self-correcting mechanisms.

There is a second point, however. Given that our characterization of the implicit Trustees' projections out to 2050 is correct, the band of possibilities represented by these outputs is rather narrow. That is, one can, as we have done, develop a range of plausible values for several of the key variables of interest that would imply a much wider range of possible outcomes in 2050 than the Trustees use in their deliberations.

This is an important issue for analysis over the long term. We often suffer from a poverty of imagination when we attempt to imagine possible future conditions. This, too, is not surprising. There is a rich context to the present that exerts a tyranny over our conception of plausible futures. By definition, those futures cannot possess the wealth of detail we unconsciously imbibe from our present surroundings. As a result, we will tend to reject or not even consider futures that lie far from what we currently perceive. And even if we, as individuals, are able to break free of these bonds, it is difficult to convey this vision to others. "Serious" analysts are well advised to stick close to the common understanding if they, in turn, wish to be taken seriously.

It is for this reason we take a more operational perspective when seeking to construct scenarios. Usually, scenarios are constructed ex ante as the Trustees have done. Different futures are imagined either, as in this case, quantitatively or as is more usual with scenario-planning exercises, qualitatively. The scenarios are then accepted (or rejected) as valid by the planning team before analysis begins. In our analysis, however, we do no pre-screening of possible future conditions. We wish, in fact, to embrace as broad a potential universe of such futures as our current state of understanding will support as being plausible. Implicit in the ranges we have selected for the thirteen unknown variables of interest are a literally countless number of scenarios -- if we take a scenario to be the playing out of a particular strategy in a future defined by the values assigned to those thirteen factors.

Why should we examine one scenario over another? To the extent we select particular scenarios for examination and analysis it is not because we have chosen those scenarios to be of interest in themselves. How would we decide? What if we are wrong? Rather, it is because those scenarios will help us understand important things about the alternatives we have and the choices we must make. Scenarios are chosen because of their ability to convey important information in our attempt to characterize the systematic successes and failures of our actual decision choices. This stands usual analytical practice on its head.

The value of this portion of the analysis is far from fully conveyed by the aggregate results in the figures. The software system permits real-time exploration of causes and effects. The output for results may be made visual and so gives users the power to look at individual or group results systematically in a wide variety of formats and from different perspectives. In this sense, the analysis provides an advance in instrumentation, permitting the implicit and the complex to be rendered explicit and comprehensible.

How Much Will We Regret Our Decisions?

Politics is the art of the possible. Similarly, in a policy analysis we are not so much interested in how well each strategy does in each of the 4,000 cases but in how well each would do compared to the other choices we have available. How much would we regret having chosen one alternative over the others?

The concept of regret is a natural complement to the concept of robustness. Once we shift from optimizing for one configuration of assumptions and instead seek robustness across many such sets, we also shift our interest from absolute to relative performance. In this context, regret compares the result from pursuing a specific strategy under a particular set of conditions to what the result would have been had we

full foreknowledge of the future state of the world and so had chosen the optimal strategy for that future. If, by some happy chance, the course we did follow is precisely the same as the one we would have chosen with perfect foresight, then our regret is zero. Any other course yields a measurable regret compared to the optimal ideal.

As an example, what if our goal is to remain comfortable and dry despite the weather? If we believe it will be sunny our optimal strategy is to not take an umbrella. If, however, we think it will rain, our optimal plan is to take an umbrella along. If the sunny day surprisingly turns into a day of rain, the regret of having followed the course we did would be getting soaked. Similarly, if the clouds suddenly parted and the skies cleared just as we committed to the umbrella strategy, our regret would be having to lug an umbrella around all day needlessly and could perhaps run to the expense of replacement if we thoughtlessly leave it in the coffee shop.

It may be at certain seasons of the year that our best information about the weather is not good enough to avoid regret one way or the other. In those cases, we might seek a strategy that would be optimal for neither of the extreme cases, if we knew for certain they would apply, but minimize our regret in each instance. In the umbrella example this might take the form of leaving the umbrella in the car. When it rains, we may get damp but only for the time it takes to retrieve the umbrella. When it is sunny there is the nuisance of additional car clutter, but we avoid greater inconvenience and minimize the chance of loss. This latter strategy might be a good candidate for a robust strategy.

Regret in 2050

We apply the same approach to the four policies in the Social Security solvency problem. We calculate the size of the balance from following the "Hold Steady" strategy, for example, in each of the 4,000 states of the world and then compare those results with what they would have been had we followed the strategy that we discover would have led to the highest trust funds' level under each of scenario. The difference is the regret of pursuing "Hold Steady" under those conditions. (For example, if "Hold Steady" leads to a $5 trillion trust funds' deficit in 2050 while the best any strategy does is a $2 trillion deficit, the regret for "Hold Steady" under these conditions would be +$2 trillion.)

The regret calculation is performed for each strategy and for each of the 4,000 alternative futures. Figure 4 shows for "Hold Steady" the number of scenarios at each level of regret. For 525 of the 4,000 futures the "Hold Steady" policy would be optimal and so the regret is zero (the first column on the left.) For another 2,400 cases, the regret would still be somewhere under $5 trillion. That is, in each of these futures had we chosen what turned out to be the optimal strategy we would have done better by

some positive amount up to $5 trillion. For the balance of the cases, however, our regret would have been over $5 trillion – in some cases very much more.

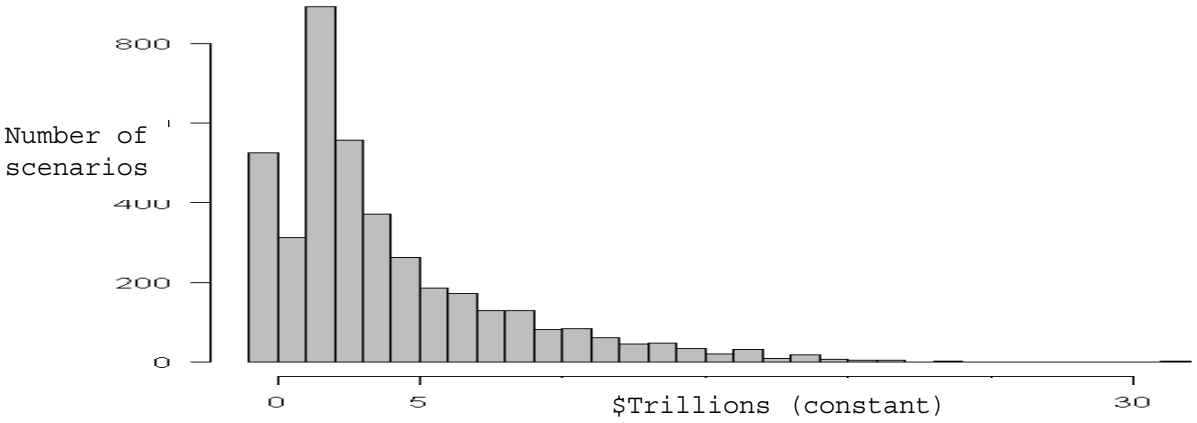

Figure 4. Regret in 2050 of "Hold Steady" Policy

Figure 5 shows a similar view for the most aggressive policy, "Adjust+Transfer". In almost 1200 scenarios this policy would be optimal, that is zero regret. Interestingly, almost the same amount shows regrets greater than $5 trillion. But the right-hand tail of large regret is not as long or thick as for "Hold Steady".

But what would cause us regret? The state of OASDI trust funds balance is one among many things we care about. We also have an interest in minimizing the tax burden necessary for maintaining the Social Security system and the indirect costs inflicted on economic growth. On this score, several of the strategies we pose above are subject to criticism. The COLA strategy seems to reduce the number of cases in which we might experience regret compared to the current policies ("Hold Steady") and to cause lower regret when it occurs. However, the simple analysis we have laid out above does not track the metric of economic well-being among Social Security recipients. Being 0.5% behind cost-of-living adjustments over a period of decades could well lead to OASDI trust funds' solvency at the cost of making old-age poverty more widespread.

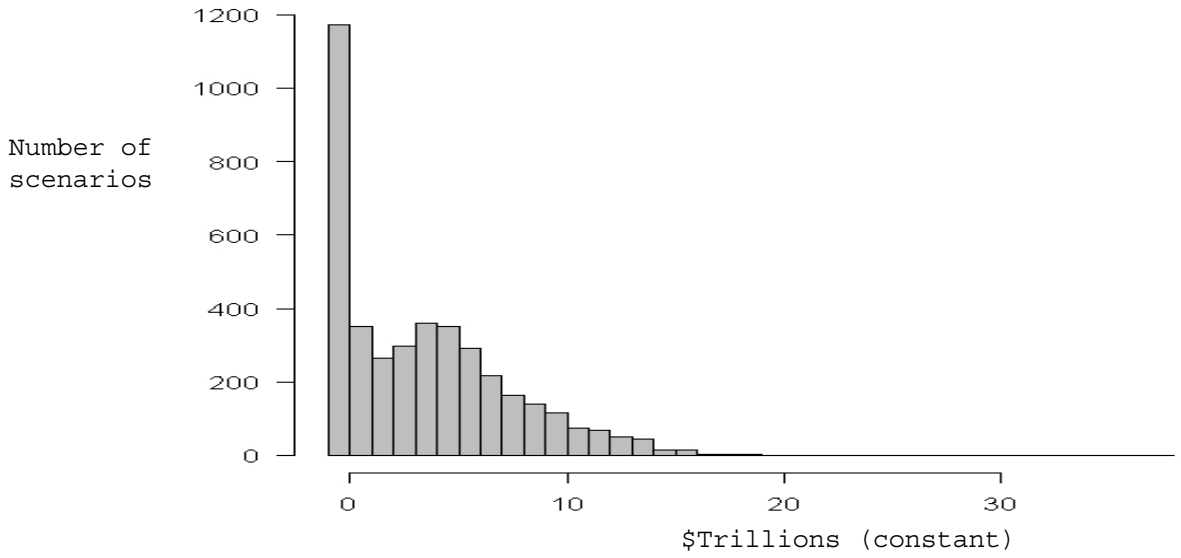

Figure 5. Regret in 2050 of "Adjust+Transfer" Policy

Similarly, while the "Adjust+Transfer" policy would seem the best bet to remove most doubts about future solvency, it does so at a potentially terrible cost. Not only would the COLA offset be double what it is in the "COLA" policy, it would also include sizable transfers from the general fund at a time when the U.S. federal budget is already running in deficit.

The model we use does not allow us to track these ancillary and secondary effects nor does it have as many internal correction mechanisms as do the models used by SSA and the Trustees. We can, however, take advantage of the peculiarity of this particular policy issue to modify the means by which we think about regret in order to reduce the degree of anomaly in the regret results.

The OASDI trust funds are not managed as a business. That is, while negative balances are avoided, there is no virtue in running too large surpluses. On the day that the funds' balances are at their lowest, a policy leading to a $1 trillion surplus may well be viewed as superior to a policy that would have left $10 trillion under the same conditions. Nine trillion of those dollars could presumably have done more good for the economy if left in the pockets of the workers who earned them.

We are not possessed of perfect foresight and the model does have its limitations, as noted above. We can, however, make a useful approximation. We know that the Trustees project a $7.92 trillion surplus in 2050 under the conditions of its Low Cost case. They do not report this figure for the Intermediate and High Cost cases because the funds go into deficit in a year before 2050 under these conditions. Therefore, we may

infer that $7.92 trillion is a level of surplus that is consistent with one or more simulations using the more sophisticated modeling available to SSA under at least one set of conditions that never goes into the red at any time between now and 2050. For this discussion, therefore, we will use this amount as a reference point.

If we examine some of the regret cases in this context it can change our perspective somewhat. For example, nearly 500 of the scenarios for the "Hold Steady" policy that exhibit non-zero regret (that is, are not the "best" strategy for those conditions) actually show positive balances of between $1 trillion and $7.92 trillion in 2050. This is because regret is a relative concept: if another policy yields a higher balance under the same conditions, then by definition one showing a lower balance, even if positive, will have non-zero regret. A further 100 regret scenarios yield positive balances in 2050 up to $1 trillion for the "Hold Steady" policy.

Similarly, the most aggressive policy, "Adjust+Transfer" raises more questions than was first apparent. Over 600 of the scenarios that show zero regret for this policy, that is where this policy yields the highest 2050 trust funds' balance, lead to balances greater than $7.92 trillion. Half of these cases show 2050 balances greater than $20 trillion. This certainly raises suspicions that it would be possible to engineer a soft landing for trust funds' solvency at less cost by choosing policies that are less aggressive. (It also illustrates how this type of analysis helps form inferences and heuristics that may then be used iteratively in refining policy options.)

Formulating a regret measure that takes into account the potential opportunity cost of too large balances requires judgment. Negative balances are clearly to be avoided if possible, but what range of positive values should we seek to target? This is at least in part a values issue that properly should be explored formally as an **"M"** quadrant matter and made a formal part of the analysis.

Choosing Among Alternative Policies

Only one of many potential compound computational experiments has been described. The analytical process is designed to be interactive and iterative. The results from one such experiment will allow people to draw inferences that will then lead to refinements for the next. All results are stored from all experiments and can be presented in accessible visualizations that permit those familiar with the policy issue, not just those with technical skills, to gain insight and provide meaningful input for the next iteration of the analysis. Through this process, individuals – both analysts and policy people – may gain a sense of what vulnerabilities matter for a variety of policy alternatives. Policy choices can be added, amalgamated or sequenced with others, and otherwise modified to

enhance robustness, in essence trying to reduce vulnerability to uncertainties as much as possible.

When a set of candidate robust strategies has been selected, they may be further tested. The same analytical system constructed for strategy construction may now be used to find scenarios that would stress a strategy's ability to meet the standards by which goodness of outcomes is measured. The ensemble of such scenarios can be made the subject of data mining techniques to describe the characteristics that lead to a stressing state of the world. That is, instead of building scenarios on an ad hoc basis, the scenarios to be examined arise naturally from the practical purpose of observing performance among our candidate strategies to determine which we should select.

It may be that a single candidate strategy dominates the other candidates when exploring across a full **"XLRM"** space and framing policies to be less vulnerable to surprise. For a sufficiently complex problem, however, it is more often the case that there remains a fundamental, irreducible uncertainty that does not permit unequivocally clear choice. Much knowledge will have been gained. Some policies will have been shown to have undesirable properties, while others will have been modified so that many of the uncertainties we feared will have had their ability to affect our goals greatly diminished. Yet, there remain policies, or variations on policies, that perform variably under different conditions.

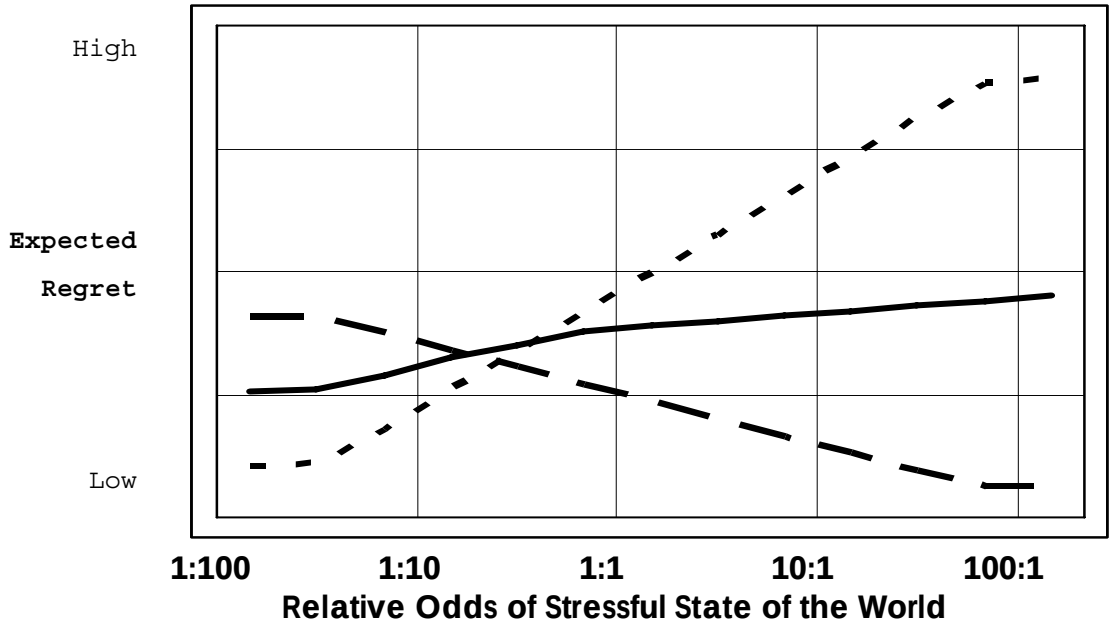

Figure 6. Robust Regions Comparison of Three Candidate Robust Strategies

When an analysis reaches such a point, it is possible to illuminate the inherent trade-offs in a manner that feeds directly into the process of

policy decision making.[20] The expected regret of each candidate policy is examined in view of varying assumptions about the likelihood that the undesirable state of the world, previously defined, would occur. Note that the remaining uncertainties are not characterized in terms of probabilities. Rather, maintaining the focus on the practical policy actions to be informed, the question is what one would have to believe were the odds of seeing particular favorable or unfavorable states of the world before you would be well advised to switch from one prospective strategic course to the other.

Figure 6 is a generalized illustration of this analytical stage. The policy whose relative regret is mapped by the dotted line performs better than the other two when the odds of a stressing future are high and so a favorable future is less likely. It performs quite poorly when these odds are low. The policy represented by the dashed line behaves the opposite although its expected regret never gets as high as that of the first policy. The policy with the straight line does not perform as well as others under either condition but has the virtue of being less affected by uncertainty. Which should be chosen? At this point we are reminded that the policy process in Congress is rooted in politics. Judgments will be made on likelihoods and costs – as they always are and always will be. But now the trade-offs and opportunities can be made clear to all.

SUPPORTING THE LEGISLATURE OF THE FUTURE – AND OF TODAY

Legislating for the future of Social Security will require rigorous analysis of policy choices in the face of explicit uncertainties. Yet, the analysis must also provide the realistic hope of being employed as part of legislative deliberations. There are several reasons for hope that this may prove to be the case for the methods the RAND project is exploring.

Improve the Ability to Observe. To be sure, the Congress suffers from institutional dysfunctions that frustrate attempts to think in the long term. But the future solvency problem is also inherently complicated (in the sense that numerous elements interact,) complex (because the results from this interaction may not be predicted,) as well as being a forum for conflicting values and interests. Even private corporations possessing more hierarchic structures, greater ability for concerted action, and a more focused "bottom line" have difficulties wrestling with problems of this character. Instruments must be provided that not only permit Congress both to generate and observe the results from "what-if?" queries that are the foundation of reasoning one's way to the future, but will also support sharing these insights within its halls.

There are several ways to carry this agenda forward. In the RAND study we begin with simple actions and characterize their performance under wide-ranging conditions. More sophisticated strategies based on

current proposals may be subjected to similar tests and examined for their responses. Modifications could then enhance these strategies against revealed dangers. This iterative approach makes practical an Occam's Razor for policy formation: use successive policy hedges to reduce or even eliminate vulnerabilities toward a portion of the uncertainties that are present.

A policy discovery approach could also be employed. Instead of beginning with fixed strategies, a more fundamental set of actions ("levers") could be defined. A suite of strategies may then be created by random or directed combinations of such levers in a manner analogous to creating the ensemble of future scenarios. Performance of these strategies under different scenarios would then be assessed. Those strategies that meet set standards of performance across a range of challenging scenarios and metrics may be dissected for their constituent elements and a tally made of these levers. Those that appear widely among the successfully robust strategies should be considered seriously for inclusion in any strategy Congress might seek to enact. Those that rarely or never appear could be dropped. In this way, the limited resources available for policy consideration may be focused on the group that falls in the middle of possible action choices.

The key to all this is the new-found ability to better understand the component elements of complex policy issues, their interaction, the options for policy, and the range of outcomes that might ensue from each possible course. The result is means for effecting dimensional collapse so that problems deeply complicated by their natures may be better comprehended by analysts and policymakers.

Shore Up Bridges Between Analysis and Policy. Legislating for the future must be a partnership, yet analysts and lawmakers are members of two separate cultures. Each has developed a vocabulary for action that works well within its sphere but often fails to translate well in the other. In particular, policymakers find the products of research answer the questions the analysts are able to answer, not necessarily the ones that are asked. Output from research becomes yet one more input; the analytical process is left in the hall when the door shuts and the true policymaking process begins.

The analytical approach outlined in this brief lowers the barrier in at least two ways. The first is that it may be conducted without first specifying assumptions about present unknowns. Indeed, the very dimensions within which the exploration occurs are explicitly defined by those uncertainties. This eliminates the need for much of the arguing over currently unknowable values that characterizes much Washington debate. The diversity of views becomes an asset in structuring more arduous testing over a wider range of assumptions. All parties are along

for the analytical ride. Analysis becomes less of an exercise in gamesmanship over what assumptions to use.

The second is that the robust decisions approach begins with the questions and choices that confront legislators. Unlike a traditional analysis that begins with framing a formal model of the system of interest and then gathers data to then yield results that will provide more or less illumination over choice, this method begins with those choices themselves and works backwards. It erects the structure necessary to discover the conditions that would cause one to reject policy 'A' in favor of policy 'B' and illuminates why this is the case and what the trade-offs entail. The analytic apparatus is framed around this core question. It has the additional value of routinely generating new choices that were not initially present. In this sense, the investigation is conducted in the same terms as is the committee room colloquy.

Robustness Arises from Adaptiveness. The robustness approach may be used to determine what constitute dangerous conditions and so identify signposts that would indicate that we are heading toward a troubling future. Hedging and shaping actions can be prepared and employed as indicated.[21]

Law and policy are often framed as if we believed conditions will not change when we know well that they will.[22] When change is then forced upon us the response is too often ad hoc and shaped by a sense of crisis and horse trading. Though inefficient, it is convenient to act this way when consensus choices are not clear. Yet there are examples of compacts that explicitly recognize that change does occur. Constitutions are sets of rules for how we will proceed as occasions arise. On this basis, groups and states feel protected in coming together for concerted action.

The robustness approach supports the framing of rules-based policy that will guide the adaptation of our actions to the requirements of the times. The strategies explored in the RAND project are of the simplest kind in specifying but one decision point, limited criteria for assessment and simple rules for adaptation. Yet, the ability to adopt a new posture in light of updated information, limited as it is, enables the possibility for response that in itself presents a more robust policy face to the deep uncertainties we confront. The robust decisions approach is a tool that has been tailored to the framing of ex ante rules for adaptive response that may be agreed upon by the various interest communities holding a stake in Social Security's future.

"Yes, But That's Not How Washington Operates". Certainly, legislating for the future is highly politicized: Politics are what Congress does. To what extent may any analytical tool, no matter how effective or well-designed for policy, hold sway in the Washington environment?

The question is important. Beyond technology and economics, it is sociology and psychology that determine how innovation is received and accepted. An historical analogy provides hope.

In the early part of the 20th century before the development of national income accounting systems, many proposals for economic policy came before Congress. Some that became law are now viewed as having been harmful to the very interests they sought to serve. The national income accounting system, a development of the 1930s and 1940s, introduced now commonplace terms -- "gross domestic product," "balance of trade", "unemployment rate", etc. – that were powerful intellectual tools for making tractable concepts previously difficult to grasp or perceive. This had two further effects. Though policy always emerges from politics, what is viewed as the acceptable trade space may change thanks to new ways of visualizing the objects of policy in concrete terms. In this way the range of economic policy proposals taken seriously as contenders for public approval has narrowed from the time before we systematized the means for inquiry and discourse. Tools from the analytical realm now place bounds on policy discussions.

Further, we will never have legislation for the future without legislators confident in being backed by constituents for the future. As previously, this is a matter for mutual education. But for this education to proceed, there must be a means for visualizing, presenting, modifying and discussing policy choices across the space of plausible futures that will support accurate presentations of trade-offs and choices in a manner that can be made readily perceptible. Until this is available and widely shared – as the once-recondite vocabulary of national income is now – legislating for the future will remain more aspiration than reality.

REFERENCES

[1] Dr. Popper is a Senior Economist at RAND. This research brief represents his personal views and does not necessarily reflect those of RAND or of its research clients or sponsors.

[2] This is the common term for the combined U.S. federal Old-Age and Survivors and Disability Insurance (OASDI) plans. The terms "OASDI" and "Social Security" will be used interchangeably.

[3] The analytical portions of this paper describe current work on a small-scale, demonstration project done in collaboration with RAND colleagues, Robert J. Lempert and Ryan Keefe.

[4] Deep uncertainty is the condition where analysts do not know, or the parties to a decision cannot agree upon, (1) the appropriate models to describe the interactions among a system's variables, (2) the probability distributions to represent uncertainty about key variables and parameters in the models, or (3) how to value the desirability of alternative outcomes

[5] The 2006 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds, 1 May 2006 (Washington: USGPO.)

[6] It is common to speak of the "Social Security trust fund" (or "funds", more properly,) but this is a problematic concept. Following the advice of a commission headed by Alan Greenspan in the 1980s, in anticipation of future shortfalls (benefits exceeding the revenues from the "pay-as-you-go" Social Security system) Social Security taxes were raised. Social Security wage tax receipts go to the Treasury to pay current OASDI benefits obligations with the surplus converted into non-negotiable special-purpose Treasury Bonds. "These [Trust Fund] balances are available to finance future benefit payments and other Trust Fund expenditures – but only in a bookkeeping sense.... They do not consist of real economic assets that can be drawn down in the future to fund benefits. Instead, they are claims on the Treasury that, when redeemed, will have to be financed by raising taxes, borrowing from the public, or reducing benefits or other expenditures. The existence of large Trust Fund balances, therefore, does not, by itself, have any impact on the Government's ability to pay benefits." *(*Office of Management and Budget*, FY 2000 Budget, Analytical Perspectives, p. 337.)* These facts allow rich grounds for differing interpretations of what actually happens when the trust funds are depleted – or even before, when outlays for benefits exceed revenues via the Treasury of OASDI tax receipts.

[7] See, "Big questions and big numbers", *Economist*, vol. 380, no. 8486, July 15, 2006, pp.67-69.

[8] A fuller discussion of the philosophy behind the robust decisions simulation approach may be found in Steven W. Popper, Robert J. Lempert and Steven C. Bankes, "Shaping the Next Hundred Years", Scientific American, April 2005. A detailed presentation of the methodology itself may be found in Robert J. Lempert, Steven W. Popper and Steven C. Bankes, Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis, RAND MR-1626 (2003.)

[9] There are several reasons why official guesses about the future don't stray far from today's norm. Among those most important for the present purpose are those related to the nature of hierarchies and bureaucracies. Herbert Simon noted the phenomenon of uncertainty absorption whereby the nuance surrounding an uncertainty

becomes lost and takes on more of a character of certainty when information necessarily becomes filtered as it is passed up the hierarchy. In this case, however the OASDI funds' Trustees may simply consider, correctly, that they already get a fairly wide spread of plausible distant futures within the range of assumptions in which they choose to conduct their analyses.

[10] These intervals were chosen after interaction with several experts in economics, demography and the social security system. The choice of range should be determined by reference to the ranges appearing both in published literature and in the statements of political figures and the representatives of the full spectrum of interest groups. The key is to embrace all views within these ranges so that all may feel their views are represented in the analysis.

[11] The policy to follow after 2020 is selected by dividing the full spectrum of scenario outcomes for all four strategies into thirds. The scenarios falling in the highest third of the range of OASDI trust funds balances either revert back to or continue the "Hold Steady" strategy. Those falling into the middle third either pursue the COLA policy (if they had been following "Hold Steady" or "COLA" since 2006) or else "Transfer" in the case of the others. The third of the cases that fall into the lowest range of 2020 balances now shift to the "Adjust+Transfer" policy in the next three decades of the simulation. The 2050 results for all 16,000 scenarios reflect the results of these adaptive responses in 2020. Again, this criterion is clearly a simple mechanical rule for the purpose of demonstration. In an actual policy analysis, the effects of choosing different rules for adaptation would be part of the exploration of the "L" quadrant of the XLRM matrix.

[12] If the year-over-year change in CPI would, under current rules, mandate a 3.5% upward benefit adjustment to correspond to the rate of inflation, the actual benefit increase would be only 3.5% - 0.5% = 3.0%.

[13] As a technical matter, it proved difficult to find revenue-side measures that could be modeled in the version of the SSASIM simulation model at our disposal. The general revenue transfer was one of the few expedients available to simulate a single revenue-side action without introducing further complications. We are aware of the fiscal, political and legal issues such a transfer policy would raise. As our focus is on demonstration of tools to support policy rather than producing an authoritative policy analysis of the trust fund solvency issue, this limitation is acceptable for the present purpose.

[14] For every $100 billion collected into the federal general fund, $500 million would be transferred immediately to the OASDI trust funds.

[15] As described by its authors, SSASIM is a social security policy simulation model under continuous development since 1994. SSASIM development has been funded by many organizations, including SSA, GAO, EBRI, AARP, IMF, Concord Coalition, and Policy Simulation Group. It actually contains two models: a macro model of aggregate program finances, the model principally utilized in the RAND analysis, and an embedded micro model of cohort individuals that can represent a wide range of policy reforms including individual accounts. Changes in assumptions and policy are accomplished by "altering values of parameters in a thoroughly documented input database, rather than having to reprogram the model's logic." (http://www.polsim.com/SSASIM.html as of 5 September 2007.

[16] The myriad combinations this would imply were sampled using a "Latin Hypercube" experimental design where we generate a set number of alternative futures spaced uniformly within the region spanned by the plausible value ranges of these thirteen variables. This, in effect, applies a uniform probability distribution across these alternative futures.

[17] This is achieved by placing SSASIM© within the CARs™ software environment. This software allows the creation, examination and modification of compound computational experiments and permits visualization of the resulting ensemble results.

[18] This is, running each of the four policies over the same 4,000 alternative futures sample: 4 x 4,000 = 16,000.

[19] The balances for the Intermediate and High Cost cases are not reported because both of these cases lead to negative balances in 2040 and 2030, respectively. Hence, the values are not calculated beyond these times. By inspection of the series of projections they do calculate, out to 2035 for the Intermediate Cost and 2025 for the High Cost cases, both sets of assumptions would probably yield a negative balance in 2050 as well. For the High Cost case the shortfall might be as much as -$10 trillion.

[20] This methodology is detailed in Robert J. Lempert, David G. Groves, Steven W. Popper and Steven C. Bankes, "A General, Analytic Method for Generating Robust Strategies and Narrative Scenarios", Management Science, April 2006.

[21] This is based on the Assumptions Based Planning approach discussed in James A. Dewar, "Improving Legislation on Long-range Issues", John Brademas Center for the Study of Congress, Research Brief-Number 4, December 2006. (http://www.wagner.nyu.edu/performance/files/Dewar.pdf).

[22] This is discussed in the context of the Kyoto Treaty in Robert J. Lempert, "Creating Constituencies for Long-Term, Radical Change" , John Brademas Center for the Study of Congress, Research Brief-Number 2, March 2007 (http://www.wagner.nyu.edu/performance/files/lempert.pdf).