

Consensus Builder:
A Place to Speak, Listen, and Be Counted

Peter Weinstein, Ph.D.
peter.weinstein@altarum.org
January 22, 2006
(White paper draft 2.0)

Consensus Builder: A Place to Speak, Listen, and Be Counted

Abstract

This paper presents a vision for an internet application that will support constructive political discourse. Consensus Builder will invite people to speak their beliefs and engage in interpretive processes that result in internal representation of statements as ontological models. Analyses of similarities and differences between statements will support listening and exchange of ideas. Semantically informed query and summarization capabilities will facilitate learning from Consensus Builder and the emergence of community leaders capable of negotiating solutions to difficult problems. Illustrations of interfaces for each aspect of Consensus Builder help the reader to understand what using Consensus Builder will be like, and to appreciate its potential contribution to society.

1 Introduction

This paper describes a new place that can potentially bring healing to many individuals and wisdom to our social policy. People will come to this place to:

- Speak about the things they know and care about
- Listen to what others have to say
- Be counted by a system that continually aggregates and publishes the beliefs of all participants.

I call this place *Consensus Builder* because through speaking and listening some people will naturally engage in creative negotiation. These exchanges will sometimes lead to solutions that can potentially provide a basis for consensus. Consensus Builder will therefore be a place that fosters constructive political discourse.

To make Consensus Builder work will require technology, experience, and knowledge from multiple disciplines including artificial intelligence, clinical psychology, sociology, and open internet communities such as Wikipedia and SourceForge. While the technical infrastructure will play a key enabling role, it will be the social organization of Consensus Builder that will really matter. In other words, the key to Consensus Builder will be the interaction between users, software agents, and human agents that will be directly or indirectly involved in many exchanges.

I am writing this white paper to introduce Consensus Builder to a variety of readers with the goal of convincing most and inspiring a few. I am a computer scientist and engineer. The discussion in this paper is focused primarily on the functionality of Consensus Builder, while bypassing issues that are truly technical. Hopefully, some readers will offer to contribute their own expertise to future papers, leading to a well-rounded description of Consensus Builder that can provide a basis for further action.

The following sections describe the Speak, Listen, and Be Counted aspects of Consensus Builder. Each section discusses motivation (why), requirements (what), and design (how) – a standard format when proposing to build software systems. A discussion section follows that addresses issues that apply to the whole concept. The paper finishes with my current thinking about next steps and conclusions.

One final preliminary comment. Consensus Builder is inherently non-partisan. In this paper I occasionally describe beliefs without a symmetrical description of the other side of the argument. This does not mean that I personally agree with the belief or that it should be the basis for social policy. Consensus Builder makes commitments on a relatively abstract level: that people have a right to express their beliefs and to be listened to; and that in some cases shared core beliefs can be nurtured to create constructive dialog.

For example, I am often struck by the similarity of emotion conveyed by proponents of both pro-life and pro-choice abortion policy. The passion on both sides seems to derive from deeply rooted concern for life. It would be hubris to guarantee that Consensus Builder can transform the harsh conflict over abortion into peaceful agreement. Most participants will come to speak and be counted – with little desire to listen. This is entirely acceptable. If even a few participants take steps towards constructive resolution of their differences, this might create movement away from fanaticism towards healing and mutual respect. At the very least, Consensus Builder will provide a mirror by which society can watch the progress of the debate and reflect on its significance.

2 Speak

This section describes the process of speaking to Consensus Builder. It seems self-evident that many people will speak their beliefs to Consensus Builder if they feel that they are being listened to and that by expressing themselves they will have at least some degree of impact. This section focuses on the nature of the experience of speaking to Consensus Builder.

We identify the following requirements:

- 1) Freedom – people should be able to speak to Consensus Builder about whatever they want, subject only to normal legal restrictions.
- 2) Attention – speakers should feel that they are being listened to.
- 3) Interpretation – the system must develop accurate models of the speaker’s intended meaning that can be used to achieve the computational behaviors described throughout this paper.
- 4) Clarity – the system should help users maximize the effectiveness of their communication by expressing their beliefs in a manner that is as coherent and well organized as possible.

The way to most thoroughly satisfy these requirements would be to hire professional staff trained to listen, help, and interpret participants’ statements. In conversation with humans, people know that they are being listened to when the listener plays close attention and demonstrates understanding with appropriate, non-judgmental and, ideally, insightful responses. Professional therapists, for example, are trained to listen well.

Unfortunately, the cost of providing a therapist for every speaker would be prohibitive, so we need to use computer technology to stretch available resources as far as possible. The state of the art in artificial intelligence, however, is very far from human intelligence. Therefore, a proposal to build a computer system that impersonates human listeners would not be credible (the ELIZA system of 1966, which parodied Rogerian therapists,

did not make any attempt to represent the meaning of user inputs and would not come close to behaving as required for Consensus Builder).

Fortunately, we can meet the requirements for freedom, attention, interpretation and clarity with technology that is clearly within our reach today. The key is to communicate clearly and honestly about the process of speaking to Consensus Builder and the expectation that participants will help the system understand their intended meaning if they want to participate fully in the consensus building process.

Thus, the process of speaking to Consensus Builder will start with participants writing what they want to say. As participants submit writing, the system starts working on interpreting the input into formal models that can be used for computation. Few participants will interact directly with the formal representations that the system uses internally, although users trained in object-oriented software design may find it interesting to do so. Instead, the system will provide a selection of interfaces that describe its current interpretation. Each interface will provide users with the ability to confirm elements of the interpretation that are accurate, and to critique and revise elements that are not accurate. Participants will choose to use those interfaces that they find most satisfying and that speak most directly to the meaning of what they are trying to express.

The interpretation process will not depend on any particular user input. We call this type of user interaction *anytime* and *anywhere* because users provide input when they want to in a manner that they select. If a user does not provide any guidance, then the expectation should be that the system will usually fail to do a good job of interpretation. To the extent that a user does engage in the process, interpretation will be more efficient and accurate. When interfaces are available that are appropriate for the user and the topic, participants should experience the process of speaking to Consensus Builder as at least satisfying, and possibly fun.

Figure 1 provides an illustration to help explain what speaking to Consensus Builder might be like. The screen is tiled with four windows. The user is free to interact with any of these at any time, and all of the standard window behaviors are available such as maximizing them to cover the full screen.

The Statement window in the upper left is where the user writes her statement. The color-coding of some of the words shows linkages to the Dialogs window below where Consensus Builder is trying to disambiguate elements of the text. For most people, writing involves a lot of erasure and editing and for the system to be trying to interpret during this process could be interruptive. Therefore, the Submit button tells the system to work on interpreting the text as revised to that moment.

The *Simple Speak* action button in the Statement window will display a version of the text that strives to use short, simple sentences that describe the essence of the statement, while minimizing subtleties such as rationales, hedging, and other nuances. The advantage of Simple Speak is that it is relatively easy for computers to understand – and sometimes relatively easy for humans to understand as well. An example of Simple Speak is provided in an appendix to this paper.

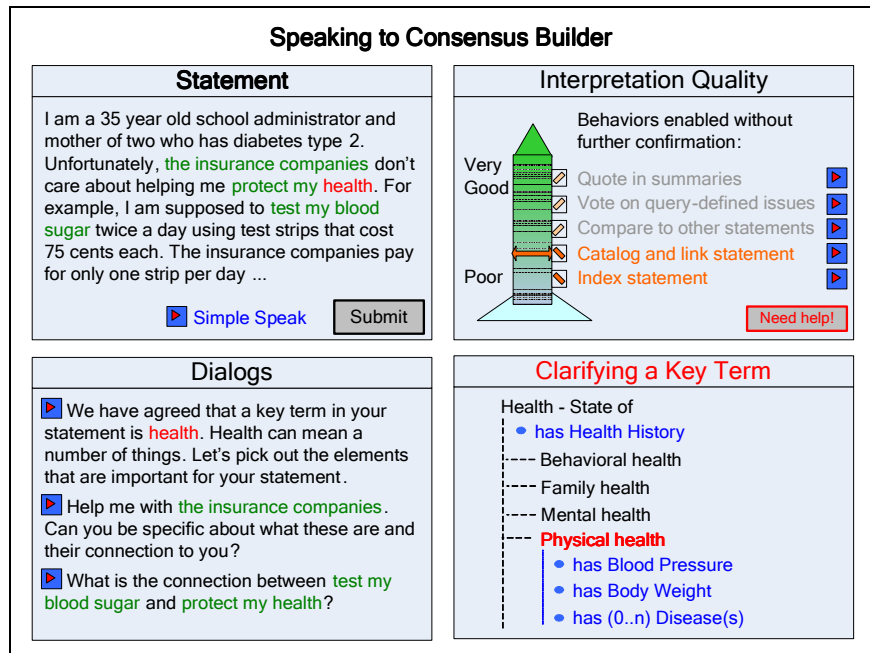


Figure 1: Speaking and interpretation

It will not be possible to automatically generate accurate Simple Speak from ordinary text (or the natural language understanding problem would be solved). It will be possible, however, for the system to generate Simple Speak from its interpretive model. Speakers will be able to edit generated Simple Speak, providing confirmation and suggestions. Simple Speak will thus be another way to engage users in a process of interpretation, complementary to the techniques discussed below.

The Dialogs window in the lower left is meant to bridge between the text in the Statement window and a variety of interpretation tools that focus on various aspects of meaning using diagrams and other devices that have clear semantics. For example, in Figure 1 earlier interaction has identified the term “health” as a key concept in the text. The system may have been clued in this direction based on the term’s place in the text, *a priori* expectations about the term for the general population, and preliminary interviews with the user. The concept of health, however, is fairly abstract. There are many dimensions to health. If the user has in mind a particular meaning and communicates this to the system then this will improve the quality of the interpretation and will lead to further useful dialog.

The Dialogs window will typically provide users with fairly substantial lists of ways to move towards improved interpretation of their statements. In the figure, the user has selected the first item and this has caused the window in the lower right to display a tool for Clarifying a Key Term. Selecting other dialog items would display different interpretation devices in the lower right.

The Clarifying a Key Term window shows an extract of a model of health in the sense of a state of being. Internally, Consensus Builder will use a type of model called *description logic ontologies* to represent meanings. Ontologies represent concepts as webs of relationships to other concepts. For example, concepts have properties such as attributes and part/whole structure. Properties are themselves represented as concepts. In Figure 1,

the concept Physical Health includes properties such as diagnosed diseases, blood pressure, etc.

A very important kind of relation between concepts is inheritance, also called “kind of”, generalization/specialization, or subsumption. For example, people are a kind of mammal, which are a kind of animal and so on. Specialized concepts inherit all of the properties of their relatively abstract parents. According to the model in Figure 1, for example, Physical Health inherits the property of having a Health History from its parent concept Health – State of.

Note that ontology as engineering practice does not imply searching for the true nature of reality as in philosophy. Rather, the idea is to construct useful models of meaning that can be applied for solving various problems. The exact content of a model will often not be critical, especially if the system is constructed in a manner that permits ontologies to differentiate and evolve as communities of discourse develop in various directions.

Like any knowledge representation, description logic has severe limits. This paper does not address the many technical tradeoffs and subtle issues. Readers familiar with the Semantic Web, however, may know that the Web Ontology Language (OWL) is a description logic (see [W3C OWL]). Consensus Builder will also use OWL to facilitate interoperability with the Semantic Web.¹

Figures 2 through 5 show other examples of interpretation devices that can be used to add precision to the system’s understanding of various elements of the speaker’s statement. These include a plan with a timeline, a data chart, a diagram of causal relations, and a map. Any diagram or other device with clear semantics can be useful. The requirement for clear semantics means that each element of the drawing has a particular meaning. For example, the arrows in Figure 2 identify temporal prerequisites (the river must be cleaned before introducing trout), while the arrows in Figure 4 identify causal relations. Consensus Builder will translate most or all of the information expressed by each interpretation device into concepts and relations in ontological models. Some devices, however, such as maps, may contain much detailed information that will not be represented explicitly by Consensus Builder. Thus, for example, the system may not be able to reason about the locations of the streams in the logging area at Sabine Falls – and this is fine unless this understanding is needed for accurate interpretation of the speaker’s statement.

¹ To understand the Semantic Web movement, imagine that we could program computers to accurately understand natural language. Internet search engines would then be even more amazingly useful! Since we do not have accurate natural language understanding in open domains, the Semantic Web proposes that web publishers provide an ontological model with every page that represents its content in a way that the computer can reason with. Major obstacles that the Semantic Web movement has not yet overcome include the difficulty that people typically experience when creating ontological models, and the need to overcome semantic heterogeneity in models created by different publishers (see Section 3 for a discussion of semantic heterogeneity).




Figure 2: A plan

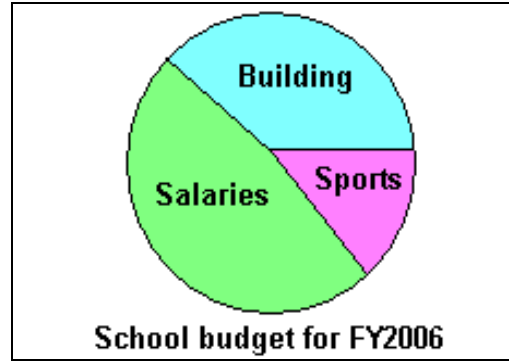


Figure 3: A data chart

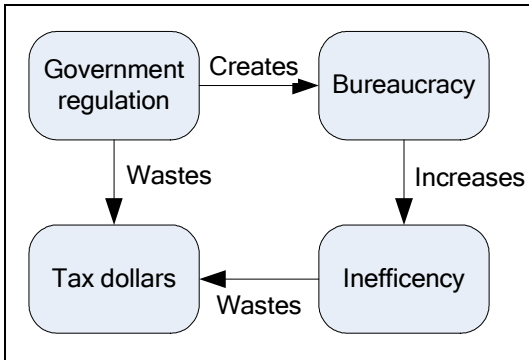


Figure 4: Causal relations

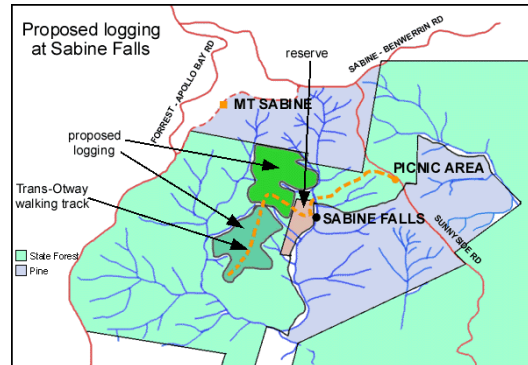


Figure 5: A map

The Interpretation Quality window in the upper right of Figure 1 reflects the system's dependence on modeling speaker statements to make Consensus Builder a place to speak, listen, and be counted. Generally, the higher the quality of interpretation achieved, the greater will be the system's willingness to behave in ways that create risk associated with possible errors. The following system behaviors, for example, will require increasing degrees of confidence in the interpretation:

- Index the speaker's statement with keyword-based technology similar to that of today's internet search engines.
- Catalog the speaker's statement with explicit hyperlinks that connect the statement and parts of the statement to those of other speakers. For example, the system (including professional staff) might develop an ontological concept for Minimize Government Bureaucracy. The speaker's statement described by Figure 4 might be classified as an instance of this concept. Furthermore, the system might automatically generate links that reference assertions within the statement. For example, other statements may focus on the connection between regulation and bureaucracy. The system would be able to link a paragraph or section of the speaker's statements to these other sources.
- Compare the speaker's statement to other statements to identify similarities and differences. This type of reasoning is discussed in Section 3 of this paper.

- Infer from the statement how the participant would want to vote on issues defined by queries. This type of reasoning is described in Section 4 of this paper.
- Extract text from the speaker's statement for inclusion in summaries that describe the state of consensus building on defined issues. This type of reasoning is also discussed in Section 4.

Each of the system behaviors listed in the Interpretation Quality window in Figure 1 is accompanied by an action button. These buttons invite speakers to provide feedback on tentative conclusions that the system infers on the basis of its current understanding. For example, the system will open new windows that show how the statement will be linked to others; or how the system believes the user would vote on various defined issues. Reasoning about these behaviors works in two directions: if the speaker confirms system conclusions this will raise the system's degree of confidence in its interpretation, while corrections to erroneous conclusions will cause revision of the interpretation.

Systems using artificial intelligence sometimes do not perform quite as well as the designers would like, and it is possible that speakers will sometimes develop a feeling of frustration during the interpretation process. The Interpretation Quality window therefore provides a button where speakers can ask for help from a human agent. The system's response to requests for help will depend on the availability of resources and other factors such as the speaker's history with the system.

A variety of human and mixed system/human interventions will be possible along a spectrum that trades off expense and quality of service. For example, if a human agent is available the system could open a text messaging window or an audio-visual connection providing an offer of unrestricted conversation. On the other side of the spectrum, this could be a structured form where speakers register complaints to be processed within a time period that reflects the current backlog, whatever that might be. The system could facilitate efficient handling of such complaints by doing its best to interpret the complaints and generating hypotheses about possible resolutions. In between these extremes, a variety of interventions could be designed to keep speakers positively engaged with the system while minimizing cost. These options could utilize professional staff, trained volunteers, or other untrained speakers – there are many possibilities.

The process of speaking to Consensus Builder described in this section will satisfy the requirements for freedom, attention, interpretation, and clarity. Speakers can state opinions on any topic and engage in anytime, anywhere interaction with the system to help it build accurate interpretive models of their statements. The process of interpretation will help speakers achieve clarity by revealing ambiguities and suggesting ways to express key ideas in simple language. Finally, Consensus Builder will convey to speakers the feeling that they are being listened to because there will in fact be listening on multiple levels: by the system, by human staff and volunteers, by other participants, and by society as a whole. The following sections describe the latter two forms of listening.

3 Listen

This section describes how Consensus Builder can encourage people to listen to other speakers, with the goal of creating constructive political discourse. Consensus Builder is not itself a tool for making decisions, but it may be consulted by decision makers and this possibility will provide motivation for people to participate in Consensus Builder as listeners.

The fundamental requirement is to create an atmosphere of mutual respect. The prevalence of negative behaviors in many existing internet forums means that Consensus Builder interactions need to be mediated in a relatively controlled way. Speakers should not be bombarded with requests, flaming must be strongly discouraged, and in general interaction must be structured in a manner that gives speakers the feeling that it is safe to express their opinions.

Perhaps a typical listening experience will start with exchanges with others with similar views, then progress towards reconciliation of more dramatic differences.

This paper focuses on the foundational technology that can be used to mediate listening, rather than the many important psychological and social issues relevant to this aspect of Consensus Builder. As with speaking, the cost of providing a human mediator for every interaction would be prohibitive. Thus, we need the computer system to carry most of the load.

We identify the following requirements (the numbering continues from the previous section):

- 5) Exchange – people should be encouraged to listen to other speakers and to engage in a constructive process of articulating their ideas while addressing the concerns of others.
- 6) Respect – Consensus Builder must maintain the feeling that it is safe to speak and that all speakers’ opinions are valued.
- 7) Comparison – To identify promising candidates for exchange and to mediate listening, Consensus Builder must be able to compare models of speaker statements. Comparisons should identify similarities and differences and quantify the degree of agreement. This capability must be achieved without placing overly restrictive constraints on the use of language: in particular, it should not restrict speaker interpretation to a single monolithic ontology.

To compare models Consensus Builder will use algorithms for graph matching. In computer science, the term “graph” means a set of vertices connected by edges. Ontological models consist of concepts connected by relations, so the models that Consensus Builder uses to represent speaker statements can be thought of as graphs. Algorithms for graph matching search for the best way to put elements of the graphs into one-to-one correspondence. For example, if concepts A1 and B1 are linked in graph 1, and concepts C2 and D2 are linked in graph 2, then if A1 is matched with C2, the algorithm will also try to match B1 with D2. In this manner, graph matching maximizes recognition of the *similarity* of the graphs. When similarity is maximized, remaining elements that cannot be matched represent the *differences* between the graphs.

Graph matching is well suited for supporting constructive discussion of political differences. A good starting point for such discussions is to explicitly recognize beliefs that are shared. Then, attention can be focused on the differences: figuring out which are important and which can be safely ignored or set aside for a later time, and looking for creative ways to reconcile the differences which are most important.

Figure 6 illustrates how Consensus Builder might support the listening process. For convenience, let us call the person who spoke her opinion in Figure 1 of the previous section Carol72. Here, Carol is comparing her statement to that of Chaim54, who also has strong beliefs about the financing of health care in America. The overall structure of interaction in Figure 6 is much like Figure 1. The screen is again split into four small windows, each of which can be resized, etc.

In the upper left window, Carol is reading Chaim’s statement. Chaim’s statement is colored green, and this color is used to identify Chaim’s beliefs in all of the listening windows.




Figure 6: Listening and analysis

The lower left window called Comparison Aspects provides a list of ways that Carol can select to compare her beliefs to those of Chaim. When speaking to Consensus Builder, both Carol and Chaim used various devices to help the system interpret their meaning. Each of these techniques modeled some aspect of their statement. Internally, Consensus Builder can use graph matching to compare all of these aspects simultaneously. The internal comparisons are too detailed and complex to be directly useful to humans, however. Instead, Consensus Builder will provide a variety of interface devices for visualizing comparisons that are each closely related to the single-perspective versions used for speaking. The Comparison Aspects window is essentially a menu. Carol has selected a causality comparison so this choice is highlighted in red.

The Causality Comparison in the lower right window overlays causal diagrams developed by Chaim (green) and Carol (blue). The color pink is used to reflect the degree that the concepts in the underlying models share definitional properties.² Thus, this visualization lets users see in an easy, natural way where the models are similar and different. In Figure 6, for example, it is clear that Carol and Chaim are more in agreement regarding their ideas about what constitutes Health Care Quality than regarding the effectiveness of American Health Insurance. Differences in the beliefs of Carol and Chaim about the effect of insurance are detailed in relations that are present in either model but not both.

The Analysis window in the upper right is currently the focus of Carol's activity. This window lists similarities and differences between Carol's and Chaim's statements. These can be inspected at different levels of detail. Carol has already spent some time ordering the differences to focus attention on what is really important. She is also adding some comments about the differences (shown in red italics) to help her prepare to write a version of her statement that addresses Chaim's beliefs.

So far our discussion of comparison in Consensus Builder has glossed over a fundamental challenge that reaches to the roots of communication and intelligence: the issue of *semantic heterogeneity*. Formally, semantic heterogeneity refers to the use of terms that are inconsistent (the same term is used to mean different things), redundant (multiple terms mean the same thing), or more generally, that have meanings that overlap in ways that can be vague and complex. Informally, the problem is that our use of language is loaded with ambiguity and there is a very large number of ways to say just about anything. As a consequence, if two speakers were to engage with Consensus Builder in processes of interpretation that were independent, the system would not be able to produce the kind of high quality comparison illustrated in Figure 6.

Fortunately, interpretation processes need not be independent at all. One way to think of interpretation is that speakers help Consensus Builder translate their statements to be expressed in terms that are defined by a globally shared ontology. The flaw in this argument, however, is that a global ontology capable of accurately expressing all speakers' views would need to contain millions or billions of definitions. To request speakers to clarify intended meanings in such a massive context would simply not work.

Indeed, to limit interpretation to any set of pre-defined terms would suppress the kind of creativity that we want to encourage. The relationship between language and thought is very close and speakers need to have the freedom to invent new terms if that helps them express their ideas.

Therefore, we are developing an approach to ontologies that we call *Living Ontologies* [Weinstein, Phelps 2006]. Living Ontologies find a productive middle ground between unrestricted use of terminology and enforced adherence to rigid standards. Figure 7 illustrates this idea informally by considering alternative ways to name the step in a business process where a person in an organization asks to purchase something. Each row

² Readers who printed this paper in black and white will have difficulty appreciating the use of color in the figure. The rectangles for American Health Insurance and Health Care Quality have three background colors, with blue on the left, pink in the middle, and green on the right.

is for a different organization. In the first column, each organization builds a model in its own way, thus producing three concepts for purchase requisition. These models will be hard to compare. In the second column, the government has imposed rigid standards that force all of the organizations to use a limited set of terms in their models. These models will be inaccurate. For example, the Navy, Army and contractor may call their processes Purchase Requisition, but in fact these processes are not quite the same. The Living Ontologies approach strives to identify maximum similarity while tolerating differences. Thus, in the third column each organization names the concept as a form of Purchase Requisition that is to some extent unique to the organization. Living Ontologies thus support models that are both accurate and comparable.

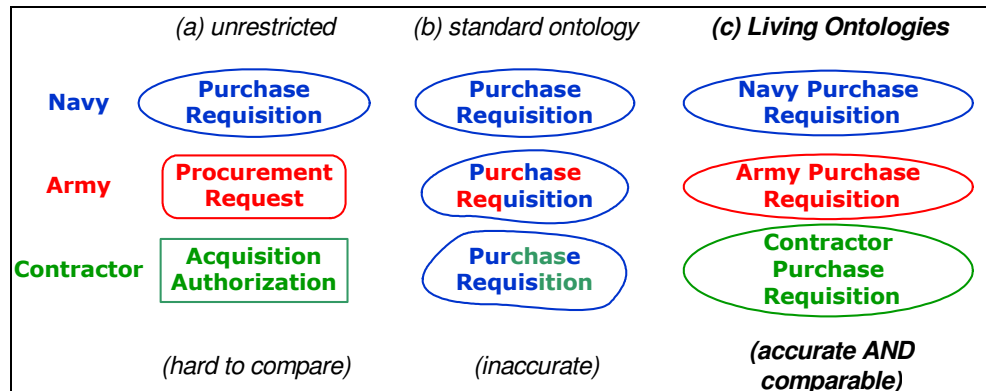


Figure 7: Three approaches to modeling and semantic heterogeneity

The key strategy for Living Ontologies is to use inheritance to separate similarities from differences. In Figure 7, for example, specialized concepts for Navy, Army, and Contractor Purchase Requisition inherit many properties from the relatively abstract concept Purchase Requisition. We call the process of maximizing inheritance *model unification*.

Figure 8 illustrates the effect of model unification on models that are specified with shared user interface devices, but that otherwise use unrestricted terminology. The concepts in these “original” models inherit some properties from core concepts that capture the semantics of the shared user interface. Model unification creates a new middle layer of generic concepts that makes similarity previously hidden in community-specific concepts explicit. The overall meaning of the unified models is not changed.

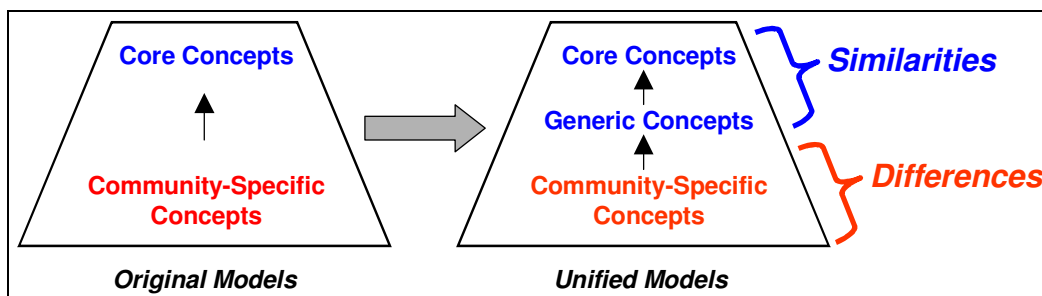


Figure 8: Before and after model unification

To accomplish model unification, we use a graph-matching algorithm that matches multiple models at the same time. This algorithm is designed in the style of

“swarm intelligence”: inspired by social insects such as ants and wasps, where organization emerges from relatively simple local interactions [Parunak 1997]. Swarming algorithms can have several benefits, including efficient identification of near-optimal solutions to hard problems, highly parallel execution, and graceful adaptation to runtime change in the problem.

Graceful adaptation is particularly valuable for creating anytime, anywhere user interaction, since each user input essentially changes the problem that the swarming algorithm is in the middle of solving. Users contribute to model unification by confirming good matches and by suggesting ways to correct errors. In general, increased user contributions improve the efficiency and quality of system performance. Users are not responsible, however, for any particular piece of input at any time.

The ant-like agents in model unification play roles defined by their places in the unified models. For example, the role of an agent associated with generic concepts is to find community-specific concepts that are best matched so as to inherit as much as possible from the generic concept. Figure 9 illustrates a matching algorithm that is analogous to the game of musical chairs. Only one concept from each original model is allowed in a match. Thus, when the organization-specific C:Prepare_RFQ concept requests to join the generic G:Create_RFQ match, the C:Create_PO concept is kicked out of the match since it is also from the original model of organization C’s business process. Agent decisions in this algorithm use estimates of concept similarity based on word co-occurrence in a corpus of documents that describe the domain, and estimates based on structural isomorphism in the original models. See [Weinstein, Phelps 2006] for details, and [Ontology Alignment] and [Ontology Matching] for research on mapping across ontologies in the context of the Semantic Web.

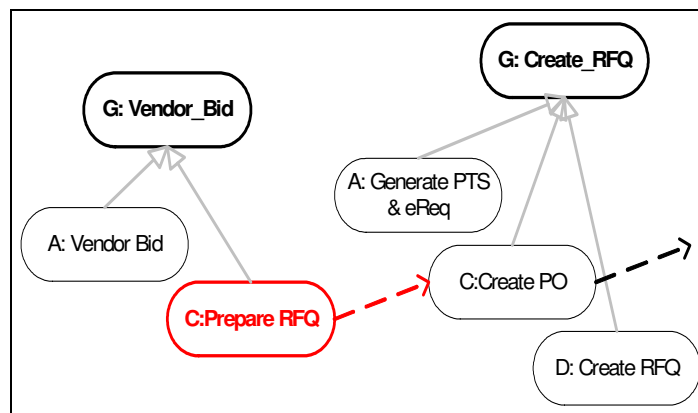


Figure 9: Concept matching as musical chairs

Consensus Builder compares models of speaker statements to support listening and exchange of ideas. This activity engages listeners using text, diagrams, and other devices focused on causality, temporal relations, etc. *Within Consensus Builder, model comparison will both require and contribute to an ongoing process of model unification.* For example, the Causality Comparison window in Figure 6 includes a concept of American generated as a match between concepts in the statements of Carol and Chaim. Hypothetically, Carol might realize that her intended meaning is different from Chaim’s partly because she considers everybody living in the United States to be American

whereas Chaim includes only citizens and legal residents. In this case, Consensus Builder would advance model unification by making Chaim's concept of American inherit from Carol's relatively inclusive version of American. Similar refinements could deal with the issue of whether people from Canada, Venezuela and so forth are also Americans.

Generally, speakers from the same community will tend to use the same ontologies. Statements of speakers from communities with little contact will tend to use ontologies that are poorly unified. In this situation, the comparisons generated by Consensus Builder will be weak.

A typical characteristic of knowledge sharing systems is that the more they are used, the greater the value created by the system since there is more knowledge to be shared. If listeners provide feedback on weak comparisons, Consensus Builder will use this information to increase the degree of model unification between their respective ontologies, leading to more useful comparisons the next time around. Opening a dialog between separate communities may require dedicated pioneers willing to delve into terminological details, but subsequent dialogs will be easier.

Thus, listening in Consensus Builder will help to build bridges between diverse communities on multiple levels. This suggests that an important role for volunteers and professional staff may be to spend time editing comparisons of statements from communities that need help to get dialog started. It may even be possible to build bridges between communities that use different languages, since there is no limitation inherent to ontologies that prohibits a mixture of English, Spanish, and other languages.

The technical infrastructure described in this section will satisfy the requirements for exchange and comparison. The requirement for respect will involve a number of issues in the psychological and social realms and will need further discussion.

4 Be Counted

This section describes how Consensus Builder can aggregate and publish the beliefs of those who have spoken. Consensus Builder will provide a stage from which individuals can address their community, nation, and the world. The greater the degree to which Consensus Builder empowers individuals, the more people will want to participate. Thus, the goal of the "Be Counted" part of Consensus Builder should be to publish the beliefs of speakers in a manner that is as persuasive and effective as possible.

We must acknowledge, however, that Consensus Builder has no claim to authority. It is tempting but misleading to describe Consensus Builder as the voice of the people – as the equivalent of a super-sophisticated, extra-flexible opinion poll. Decision makers in representative democracies are very sensitive to polls. It will be wonderful if they are also very sensitive to Consensus Builder. Scientifically conducted polls, however, estimate the opinions of the full population within statistically determined confidence intervals. The mathematics of polls is based on randomly selected samples. Consensus Builder violates the requirement for random samples by encouraging people to speak about the issues that matter most to them, and by the self-selection of people who decide to participate. While it might be possible to add scientific polling to Consensus Builder, this might be awkward.

Rather, the persuasive power of Consensus Builder will depend on the quality of its results. The underlying assumption is a belief in human problem solving under conditions of mutual respect and the exchange of ideas. The design of the “Be Counted” part of Consensus Builder should focus on how to enrich and facilitate speaking and listening: in other words, the problem solving process. Thus, Consensus Builder should be a tool for learning.

We identify the following requirements:

- 8) Transparency – Since it is impossible to remove subjectivity and bias from the process of summarizing complex information, Consensus Builder should instead strive to make analytic assumptions visible.
- 9) Community – Consensus Builder should foster the development of leaders: people who become recognized as speakers for their cause, who are thus well positioned to engage in a process of listening and negotiation.

Figure 10 illustrates some of the capabilities that might be developed to support learning from Consensus Builder. The counts displayed have no basis in data.

The Query window in the upper left shows a query specified as a short sentence in natural language. Relatively complex queries will also be possible, although these might require interpretation as when speaking to Consensus Builder. We might anticipate that many users of Consensus Builder will first become comfortable with the process of interpretation when submitting queries, which are typically much less complex than statements.

Queries specified as phrases or sentences carry substantially more information than the lists of words used by internet search engines today. The information content of these queries will be commensurate with the internal representation of speaker statements as ontological models. The resulting accuracy of the responses will be substantially improved compared to today’s search engines, with respect to both precision and recall. This improvement and the features described below are straightforward benefits of organizing information with description logic, much as is promised for the Semantic Web in general [Daconta et al. 2003].

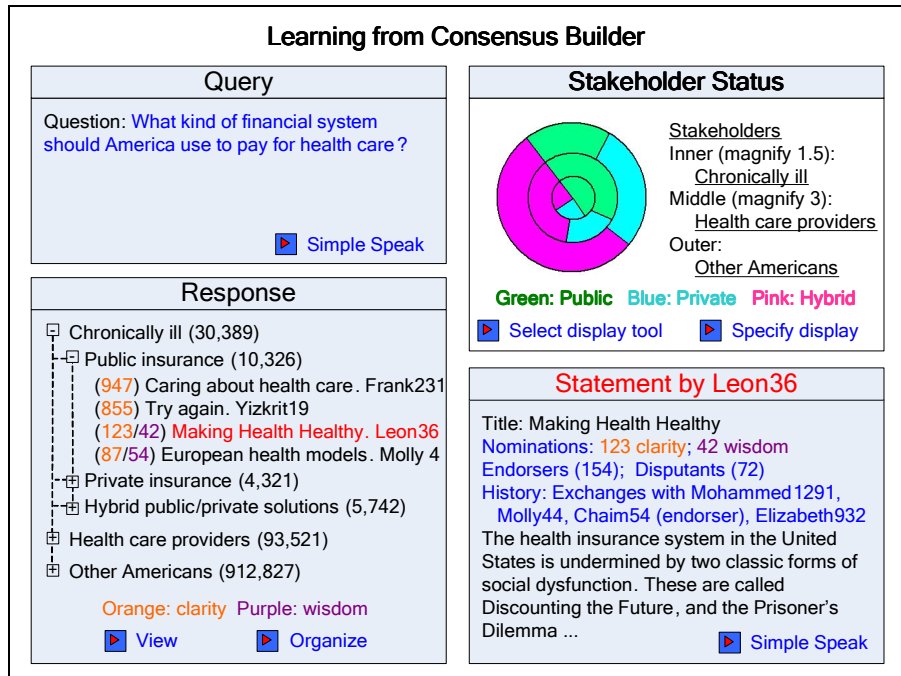


Figure 10: Learning from Consensus Builder

The Response window in the lower left of Figure 10 shows how Consensus Builder can use its understanding of speaker statements to organize query responses in flexible and meaningful ways. In the figure, the user has previously used the Organize action to define a categorization of the responses according to metadata provided by the authors (distinguishing Americans who are chronically ill, health care providers, and other Americans), and according to the substance of the statements (whether a public, private, or hybrid financial system is preferred). The numbers in parentheses to the right of each category is the total number of statements that are members of the category. To categorize statements according to their content is essentially to vote on the issue defined by the query – a form of reasoning listed in the Interpretation Quality window of Figure 1 in Section 2.

The numbers in orange and purple associated with statements in the Response window are measures of recognition awarded to these statements by other participants. The design of these measures will be very important for satisfying the Community requirement. For example, participants may be granted the right to nominate a limited number of statements for recognition of various aspects of quality. The statements listed in Figure 10 are sorted by the orange measure. Other sort orders could include relevance to the query, counts of links to the statement (a strategy utilized by Google), date of publication, and so on. Selecting a statement and clicking the View action causes display of the statement in the Statement window in the lower right of Figure 10.

The Statement window in Figure 10 is much as we’ve seen before, except that metadata describing the statement is displayed along with its text. Here the orange and purple measures are labeled as “clarity” and “wisdom”, but the definition of these or other measures of recognition is yet to be determined. Also shown is the number of participants expressing strong agreement with the statement (Endorsers) and strong disagreement

(Disputants). Statements that achieve high recognition for quality, many endorsers and few or zero disputants may be indicative of consensus.

The Stakeholder Status window includes a visual summary of beliefs on the issue defined in the Query window, structured according to the organization defined in the Response window. Here, the chronically ill are displayed as the inner circle of stakeholders, while health care providers are the middle ring and other Americans the outer. The area within each ring is colored according to the position taken by the statement regarding the issue defined by the query, in proportion to the number of statements in each category. This visualization device highlights discrepancies in frequency of beliefs between the layers of stakeholders.

The analyst uses the Specify Display action to set aspects of the visualization that complement the organization of the response. These display decisions will often require subjective judgments. For example, it is not objectively certain that the chronically ill should constitute the inner circle of stakeholders. Furthermore, members of the inner circle are relatively few in number but have a relatively large stake in the outcome. If the area of each ring were proportional to the number of statements, the visual impact of the inner circle would be too small. Therefore, the analyst has specified that the area covered by the inner circle be magnified by a factor of 1.5 relative to the middle ring, which is itself magnified by three relative to the outer circle. Identification of these scaling factors in the legend is an example of the transparency desired for Consensus Builder displays. Transparency makes subjectivity visible, helping readers to evaluate Consensus Builder results critically as is appropriate for any media content.

The Select Display Tool action will provide a menu of devices for visualizing various kinds of analyses. One way to use these results will be to copy and paste them into new statements spoken to Consensus Builder. Thus understanding created using Consensus Builder will enrich further discussion.

Many other capabilities to facilitate learning from Consensus Builder could be developed in addition to those illustrated in Figure 10. Some ideas include:

- Textual summaries of the state of discourse, including quotes and drill-down into the full statements.
- Visualizations of similarity and differences across communities of belief. Visualization of shared layers of belief (such as the concern for life shown by both sides of the abortion debate) might be powerful in an artistic sense, although not necessarily interesting from an analytical point of view.
- Analyses of the structure of discourse as it develops over time. For example, one might imagine a crystallization of coherent positions that then supports a process of negotiation and reconciliation. Understanding patterns in the development of consensus or polarization would certainly be of great interest to social scientists, and to stakeholders in particular issues.
- Identification of unusual statements – outliers that do not fit well into any category, that may contain seeds for new and potentially productive ways to look at problems.

This section describes how Consensus Builder can meet the requirements for transparency and community. Consensus Builder can become a powerful tool for learning about issues as experienced by all stakeholders. In the process of helping participants learn about the beliefs of others, Consensus Builder will fulfill its promise to all speakers that they will be counted.

5 Discussion

This section briefly discusses two issues that pertain to all aspects of Consensus Builder.

5.1 Identification

The world is a dangerous place and people of all ages are commonly advised against revealing personal information on the internet. Most people speaking to Consensus Builder should therefore identify themselves with screen names rather than legal names.

On the other hand, to provide full anonymity would make Consensus Builder vulnerable to distortions by those willing to create multiple screen names. Therefore, it seems that screen names must be registered in association with information that can be used for authentication, and to enforce a rule of one screen name per person.

Some people may prefer to speak to Consensus Builder using their true name, especially if they are already public figures.

Metadata describing speaker attributes such as nationality, ethnicity, gender, profession, etc. will be invaluable to the consensus building process, but should not be specified to a level of detail adequate to establish identity except for those using their true names. The system will need to provide strong security to ensure that personal identities remain private.

5.2 Trust

The overall thrust of Consensus Builder is strongly democratic since its fundamental objective is to provide a way for people to participate in public political discourse. There is a possibility, however, that the system could be manipulated.

To create and protect the trust essential for Consensus Builder to succeed, Consensus Builder must be:

- Non-partisan
- Non-profit
- Open source.

Open source means that the computer code that implements Consensus Builder should be available to the public. Development should be conducted in a collaborative manner under the auspices of SourceForge, Apache, or another open source development community.

6 Next Steps

Here are some things that need to be done to move forward with Consensus Builder.

- Fill holes in the research and broaden its scope.

For example, I believe this paper is too simplistic in its thinking about social issues. To illustrate, two recent books point out issues that need further thought. In The Wisdom of Crowds, James Surowiecki describes the advantages of inclusion, diversity, and decentralization for decision making [Surowiecki 2004]. But, there is a delicate balance to be struck between the needs for independence of thought and aggregation of opinions. A closer analysis of these issues will result in improvements to the Listen and Be Counted aspects of Consensus Builder.

In The Toyota Way, Jeffrey Liker describes the advantages of building consensus – called *nemawashi* – for corporate decision making [Liker 2004]. The extra time required to make decisions via *nemawashi* is more than offset by the ability to rapidly implement decisions once they are made. Not surprisingly, the principles of *nemawashi* are entirely consistent with Consensus Builder. But, one of the requirements of *nemawashi* is “set-based concurrent engineering” – a thorough consideration of alternatives. This is something we have not yet accomplished for Consensus Builder.

There are also technical difficulties that need to be worked through. For example, description logic is insufficiently expressive to capture the full meaning of second order relations such as “believe”, “should”, and indicators of uncertainty such as “may”. This will result in a certain amount of complexity in the computer code for interpretation and comparisons. In comparison to the Semantic Web in general, however, Consensus Builder is a relatively modest endeavor from a technical point of view. In Consensus Builder, ontological models are generated by a single system with core ontologies shared by all. Also, Consensus Builder provides a well defined, relatively tidy task context compared to the Semantic Web, which seeks to model most or all human knowledge.

- Establish a home on the internet including an initial collaboration infrastructure.

The collaboration infrastructure should be organized to support set-based engineering, using inheritance to organize alternative approaches and composition to articulate structure and behavior. The work at this point should move beyond narrative to include efforts in software and organizational design.

It will probably be helpful to start with an existing tool. For example, the collaborative web-editing tool “twiki” is available at no cost to the public. It may also be advantageous to join an existing collaborative community. Wikipedia, for example, is a collaborative online encyclopedia that also provides a home for projects outside the scope of writing the encyclopedia.

- Seek an initial sponsoring organization.

Consensus Builder is too big a project to be executed by an individual, a small team or even a small organization. To start we need to find at least one sponsoring organization.

- Seek publicity and funding to increase participation in the project.

It seems to me that internet-assisted word of mouth and other relatively low-key techniques for seeking publicity would be most consistent with the participative philosophy of Consensus Builder.

- Build a community of human agents with deep understanding of the psychological, social, and technical aspects of consensus building.

Human agents will play an important role in Consensus Builder, filling the gap between the level of intelligence required and the kind of intelligence that can currently be achieved with computers. It will be vital that a community of human agents deeply committed to the core principles of consensus building exists before opening the tool to the public.

7 Conclusions

This paper has described a vision for an internet application that can organize and facilitate political dialog involving very large numbers of participants. The motivating premise is that high quality solutions will emerge from a process where knowledge is shared and discussed in an atmosphere of creative negotiation. Consensus Builder itself has no authority to make decisions directly (which may help to reduce the vitriol that often accompanies political debate). Rather, Consensus Builder is a tool for generating recommendations whose persuasiveness is determined by the quality and relevance of their content.

Health care in the United States, including the issues mentioned by the examples in this paper, is particularly attractive as an initial focus for Consensus Builder. This is because:

- Health care affects everybody.
- There is a clear need for change. Escalating costs are putting pressure on businesses and financial failure is looming for major public programs such as Medicaid. The 1992 effort to reform health care failed, and the new prescription drug benefit seems almost ridiculously complex. Many Americans may agree that our government seems to need help.
- Health as a topic brings out the cooperative and caring nature of people. Since health is intimately entwined with death, it touches the deepest levels of our shared humanity.

There is no guarantee that a nationwide discussion on health will produce consensus, or even a small set of clearly differentiated ideas. It would be somewhat discouraging if Consensus Builder yields an inconclusive mass of ideas and opinions (at least it would be a well-organized mass thanks to the classification capability of description logic). This result might indicate the need for a different mediation strategy.

Even in this scenario, however, one would anticipate many benefits that would come from the creative energy generated when large numbers of people become involved in social issues. Opportunities for improvement exist at all levels of organization. And change often creates further opportunities for improvement.

To create a national discussion about health will require the involvement of hospitals, insurance companies, state governments, federal agencies, educational institutions, hospices, and other organizations. It so happens that the Altarum Institute, a non-profit organization that is currently my employer, is now undergoing a reconfiguration with the goal of becoming a leader of national health transformation. Altarum may be well

positioned to catalyze the creation of such a coalition. It seems to me that the vision of Consensus Builder – as an infrastructure that can make a national discussion about health feasible and productive – can immediately bring value to Altarum in its bid to achieve leadership.

Feedback from readers including questions, critique, suggestions, and other contributions will be most appreciated.

8 References

- [Daconta et al. 2003] Daconta, M. C., Obrst, L. J., Smth, K. T. (2003). The Semantic Web. Wiley Publishing, Indianapolis, IN, USA.
- [Liker 2004] Liker, Jeffrey K. 2004. The Toyota Way. McGraw-Hill, New York, New York, USA.
- [Ontology Alignment] The Ontology Alignment Source. 2003.
<http://www.atl.lmco.com/projects/ontology/>.
- [Ontology Matching] Ontology Matching. 2005. <http://www.ontologymatching.org/>.
- [Parunak 1997] Parunak, H. Van. 1997. "Go to the Ant": Engineering Principles from Natural Agent Systems (1/97). *Annals of Operations Research* 75 (1997) 69-101.
- [Surowiecki 2004] Surowiecki, James. 2004. The Wisdom of Crowds. Doubleday, New York, New York, USA.
- [W3C OWL] OWL Web Ontology Language Overview. 2006.
<http://www.w3.org/TR/2004/REC-owl-features-20040210/>.
- [Weinstein, Phelps 2006] Weinstein, Peter, Tom Phelps. 2006. Business Process Interoperability with Living Ontologies. Altarum Institute Technical Report. Ann Arbor, Michigan, USA.

9 Appendix A – An Example of Simple Speak

This appendix includes an example of a statement translated into Simple Speak. Simple Speak may play an important role in Consensus Builder by helping to bridge the gap between how people use language and the computer's way of representing knowledge (see Section 2 of the white paper). As the ability of Consensus Builder to understand natural language improves over time, Simple Speak will become less and less needed. This will also apply to other techniques for interpretation.

The features of Simple Speak include:

- Many Simple Speak sentences are phrases in the form of < subject relation object >. Other Simple Speak sentences link a small number of phrases together.
- Entities are referenced in a repetitive manner rather than using pronouns or other substitutions.
- The use and scope of negation (not) and disjunction (or) is minimized.
- Explicit statement of some relations implicit in natural language.

Speaker statement

I am a 35 year old school administrator and mother of two who has diabetes type 2. Unfortunately, the insurance companies don't care about helping me protect my health. For example, I am supposed to test my blood sugar twice a day using test strips that cost 75 cents each. The insurance companies pay for only one strip per day. 75 cents a day may not seem like much but for a person in my situation it adds up. Besides, testing my blood regularly means that I am much less likely to need to amputate a foot or do some other awful and expensive procedure in ten or twenty years. So it makes sense to pay the smaller amount now to help me stay healthy and productive!

Simple Speak

I am 35 years old. I am a school administrator. I am a mother. I have two children.

I have diabetes type 2. People with diabetes type 2 do tests. The tests measure the level of sugar in the blood. People with diabetes type 2 should do two tests each day.

Insurance companies do not care about my health. Insurance companies pay for tests. Insurance companies pay for one test each day. I pay for one test each day. Paying for one test each day is hard for me.

The tests protect my health. I need good health to be productive. People with diabetes type 2 often need amputations. I am afraid I may need an amputation. The tests protect me from needing an amputation.

The tests are not expensive. Amputation is expensive. Insurance companies pay for amputations. So insurance companies should pay for two tests each day!