


# Collective Causality: Building Solution Architectures With a Crowd

by  
**Carolyn J. Fu**

M.S. Mechanical Engineering (2010)  
Stanford University

B.S. Mechanical Engineering (2009)  
University of Pennsylvania



ARCHIVES

Submitted to the System Design and Management Program  
in Partial Fulfillment of the Requirements for the Degree of

**Master of Science in Engineering and Management**

at the  
Massachusetts Institute of Technology

June 2017

© 2017 Carolyn J. Fu  
All rights reserved

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

**Signature redacted**

Signature of Author \_\_\_\_\_

Carolyn J. Fu

System Design and Management Program

May 12, 2017

**Signature redacted**

Certified by \_\_\_\_\_

Professor Thomas W. Malone

Thesis Supervisor

Patrick J. McGovern (1959) Professor of Management

**Signature redacted**

Accepted by \_\_\_\_\_

Joan Rubin

Executive Director, System Design & Management Program



77 Massachusetts Avenue  
Cambridge, MA 02139  
<http://libraries.mit.edu/ask>

## **DISCLAIMER NOTICE**

Due to the condition of the original material, there are unavoidable flaws in this reproduction. We have made every effort possible to provide you with the best copy available.

Thank you.

**The images contained in this document are of the best quality available.**

*This page is intentionally left blank.*

# **Collective Causality: Building Solution Architectures With a Crowd**

By

Carolyn J. Fu

Submitted to the System Design and Management Program on May 12, 2017 in Partial Fulfillment of the Requirements for the Degree of Master of Science in Engineering and Management

## **Abstract**

Traditional open innovation has operated on the assumption that by casting a wide net into the crowd, the likelihood of obtaining a desirable solution to a problem increases, due to the greater range of potential solutions that is obtained. This is typically implemented using a competitive format, where the best ideas are selected from a crowd, and the rest are discarded.

Unfortunately, the drawback of such a format is that it fails to make use of the efforts behind discarded ideas. Each of these ideas represents a great deal of cognitive effort that has gone towards understanding and solving a problem, and discarding them sacrifices potentially useful insights that might be derived from ultimately unworkable solutions.

This thesis explores how a more effective form of collective intelligence might be obtained – one where the half-baked solutions of many participants might be combined to produce something more effective than one participant's fully baked solution that is selected through competition.

The specific format of a collaborative causal map is explored, where individuals can each contribute causes and causal links to an overall causal web, building an ever richer architecture of potential solutions (and their sub-solutions) to an overall problem. The goal is to integrate individuals' contributions such that they accumulate to an overall cohesive solution that is better than what any individual could have developed.

A series of pilots are conducted to understand the group dynamics in both offline and online collaboration, and determine those factors that are material to the success of an online collaborative causal map. Such factors include how the question is framed, how users attend to others' contributions, or how users' contributions can be curated. These factors are ultimately incorporated into a prototype collaborative causal mapping website, which is developed for public use.

Thesis Supervisor: Professor Thomas W. Malone

Title: Patrick J. McGovern (1959) Professor of Management



*This page is intentionally left blank*

## **Acknowledgements**

Theses appear through the incredibly generous support of many people. This work is the product of professors, friends and family who have collectively brought these thoughts into existence, and guided their articulation.

Leading the charge in this collaboration has been Professor Thomas Malone, whose imagination was everything I hoped for and more, in the leap that pulled me across the sea to MIT. His guidance, ranging from the technical to philosophical, is responsible for just about every causal link that holds the architecture of this thesis together.

Through his seminar, Professor Ray Reagans became an irreplaceable chaperone through organizational theory. He guided my exploration and exploitation across the landscape of how groups become more than the sum of their parts, and opened the door to a peak of my very own.

The greatest gift this thesis has given is the opportunity to not let it go. I am incredibly grateful to the faculty of Sloan who have provided me the liberty to linger, and will ensure that I never run out of questions to answer, as I find my way to the most important ones.

The faculty, staff and friends in SDM have also been instrumental in this journey. Professor Bryan Moser laid the foundation for understanding teamwork in complex projects; Jolly, JD, Erdem and Dave gave their invaluable comments on this project and tested many a bug out of its unwieldy prototypes.

My family and friends around the world have removed so many of the negative externalities that I desperately feared this journey would incur. I can never be grateful enough for every act or thought that has made the world feel so small.

Lastly, I have the greatest gratitude for a significant individual who came along for the journey from the very first day, and waded through its muddy waters and rode its rollercoaster peaks. You have given me a confidence and imagination I never dared dream of.

*This page is intentionally left blank.*

**Table of Contents**

- 1 Introduction ..... 10**
  - 1.1 How Open Innovation is Currently Structured ..... 10**
  - 1.2 Moving Away From Natural Selection..... 10**
  - 1.3 Aiming Towards A Collective *Solution Architecture* Platform ..... 10**
  - 1.4 Structure of This Thesis ..... 11**
- 2 The Architecture of Solutions ..... 12**
  - 2.1 Solutions Versus Solution Architectures ..... 12**
  - 2.2 Leveraging Open Innovation to Build Solution Architectures ..... 13**
  - 2.3 How Solution Architectures Are Integrated ..... 14**
  - 2.4 Types of Solution Architectures ..... 16**
  - 2.5 Platform Goals..... 17**
- 3 Discovering the Dynamics of Collective Causality ..... 18**
  - 3.1 In-Person Collaboration ..... 18**
    - 3.1.1 Analysis ..... 20
    - 3.1.2 Takeaways ..... 20
  - 3.2 Online Collaboration – Freeform versus Guided Map..... 22**
    - 3.2.1 Freeform Collaboration ..... 22
    - 3.2.2 Freeform Alternative ..... 23
    - 3.2.3 Guided Map Collaboration..... 25
    - 3.2.4 Takeaways ..... 30
  - 3.3 Asking the Right Question ..... 30**
    - 3.3.1 Question Framing..... 30
    - 3.3.2 Anchoring Bias..... 33
    - 3.3.3 Availability Bias ..... 36
    - 3.3.4 Takeaways ..... 38
  - 3.4 Curation ..... 39**
    - 3.4.1 Likelihood as a curator ..... 39
    - 3.4.2 Takeaways ..... 42
- 4 Collective Causality Platform Design..... 44**
  - 4.1 Summarized Takeaways..... 44**
  - 4.2 Platform Development..... 46**
  - 4.3 Implementation of Design Features ..... 47**
    - 4.3.1 Map Visibility: Random Versus Whole..... 47
    - 4.3.2 Seed Nodes ..... 47
    - 4.3.3 Simple Outcome and Fill-In-The-Blanks ..... 48
    - 4.3.4 Control Character Limit ..... 49
    - 4.3.5 Personal Refinement..... 50
    - 4.3.6 Likelihood Curation..... 51
  - 4.4 Platform Usage ..... 51**
- 5 Conclusion ..... 53**
- References..... 54**

## List of Figures

Figure 1. Eventual causal map output .....	19
Figure 2. Group undergoing brainstorming discussion to develop causal map.....	19
Figure 3. Prompt for users to contribute to a communal proposal.....	22
Figure 4. Sample output #1 of freeform collaboration, where ideas are listed separately .....	23
Figure 5. Sample output #2 of freeform collaboration, where ideas are in the same prose, but ideas remain independent of one another .....	23
Figure 6. Example prompt to improve on a given prose.....	24
Figure 7. Subsequent improvements upon the proposal by each user.....	25
Figure 8. Basic causal mapping platform developed for the pilot .....	26
Figure 9. Users' contributions to a collaborative causal map .....	27
Figure 10. Lower node cannot be understood independent of upper node .....	27
Figure 11. Node below actually includes two integrated concepts. In addition, it cannot be understood independent of node above.....	28
Figure 12. Crowd-generated causal map without minimum character limits .....	29
Figure 13. Prompts given to Turkers under "simple cause" condition.....	31
Figure 14. Prompts given to Turkers under OPM condition .....	32
Figure 15. Prompts given to Turkers under "Outcome" condition.....	32
Figure 16. Turkers asked to evaluate if Event A and Event B were related.....	33
Figure 17. Varying visibility of map to test anchoring bias.....	34
Figure 18. Turkers asked to evaluate creativity and practicality of a contribution.....	35
Figure 19. Responses from condition where user is anchored by other contributions.....	35
Figure 20. Responses from condition where no other contributions are shown .....	36
Figure 21. Task allowing Turkers to select their own prompt .....	37
Figure 22. Frequency of cause prompts selected by users, based on order of appearance..	37
Figure 23. Frequency of cause prompts selected by users, based on order of appearance..	38
Figure 24. Cause generation task given iteratively to Turkers, where each “\${event}” was a contribution from a previous round’s result .....	40
Figure 25. Subset of the map produced by the iterative causal mapping exercise .....	40
Figure 26. One of the full 5-level deep chains produced (from political influence to CO2 level reduction) .....	40
Figure 27. Likelihood evaluation of cause effect pairs given to Turkers .....	41
Figure 28. Chain of events presented to Turkers.....	41
Figure 29. Variation of quality as a function of average likelihood.....	42
Figure 30. Variation of quality as a function of likelihood standard deviation.....	43
Figure 31. Prototype platform created to test causal mapping in the context of the November 2016 elections .....	46
Figure 32. Different opacity of elements of different likelihood in the map .....	47
Figure 33. Seeded policies and outcomes created before publication of the website.....	48
Figure 34. Prompt to contribute an outcome on mouseover of an existing event .....	49
Figure 35. Prompt to fill in an outcome.....	49
Figure 36. Error alert when submission is not between 10 and 200 character limit.....	50
Figure 37. Editing interface .....	50
Figure 38. Deleting interface.....	50

Figure 39. Users can vote on the likelihood of each relationship .....	51
--	----

**List of Tables**

Table 1. Summary of Pilot #1 parameters.....	18
Table 2. Takeaways from in-person pilot .....	21
Table 3. Summary of Pilot #2 parameters.....	22
Table 4. Summary of Pilot #3 parameters.....	24
Table 5. Summary of Pilot #4 parameters.....	25
Table 6. Takeaways from online pilots.....	30
Table 7. Summary of Pilot #5a parameters .....	31
Table 8. Summary of Pilot #5b parameters.....	32
Table 9. Summary of Pilot #6a parameters .....	34
Table 10. Summary of Pilot #6b parameters .....	34
Table 11. Summary of Pilot #7a parameters.....	37
Table 12. Takeaways on how to solicit contributions from users.....	39
Table 13. Summary of Pilot #8a parameters.....	40
Table 14. Summary of Pilot #8b parameters .....	41
Table 15. Summary of Pilot #8c parameters.....	41
Table 16. Takeaways on how to implement curation .....	43
Table 17. Summary of takeaways from pilots.....	44
Table 18. Summary of features considered for prototype .....	45
Table 19. Summary of user interactions with prototype platform .....	51

# 1 Introduction

This thesis will explore how the ideas of multiple online participants can be coordinated so as to develop better solution architectures than what individual participants are able to achieve alone. Specifically, this will be examined in the context of crowds developing causal chains together through an online platform.

The thesis will explore both (1) design parameters that might impact the usability and usefulness of such a platform, and (2) the implementation of an actual platform based on these parameters.

## 1.1 How Open Innovation is Currently Structured

Traditional open innovation has operated on the assumption that by casting a wide net into the crowd, the likelihood of obtaining a desirable solution to a given problem increases – both because a larger variance of potential solutions is achieved, and also because individuals with diverse experiences can often provide novel approaches to tackling longstanding challenges. In such open innovation competitions, organizations typically pose a question to a crowd, and select the best answer amongst the responses from individuals or teams, and discard the rest.

## 1.2 Moving Away From Natural Selection

Consequently, this process of natural selection fails to make use of the efforts behind these discarded ideas. Much of the cognitive effort that has gone towards understanding and solving the problem is wasted, potentially sacrificing what could have been valuable insights within what were ultimately unworkable solutions. One can perhaps imagine a more effective format of collective intelligence, where the half-baked solutions of many participants might perhaps be combined to produce something even more effective than the fully baked solution that was ultimately selected.

Furthermore, and perhaps more pertinently, the natural selection approach to open innovation may in fact no longer be viable for the types of complex problems we aim to solve today. As problems begin to cut across an increasing number of domains, societies, institutions, and so on, it is unlikely that one individual or even team, whether inside or outside an organization, will have the whole solution to a problem.

## 1.3 Aiming Towards A Collective *Solution Architecture* Platform

It thus behooves us to find a way that we can more effectively collect the diverse insights of individuals across many areas, so as to grow a solution that could not have been developed by any one individual or team alone. We can envision a platform that will enable large numbers of individuals to contribute individual parts of a solution, in a way that can be effectively amassed online. We aim to use the crowd to contribute to a collective *solution architecture*, as opposed to having them provide *solution alternatives* that are selected over.

#### **1.4 Structure of This Thesis**

Developing such a solution is a two part problem – first, understanding what it is that enables effective group innovation in general, offline contexts; and secondly, understanding how this can be achieved using the levers available in the design of existing collective intelligence platforms. This thesis will tackle both parts of this question.

In Chapter 2, the concept of solution architectures will be expounded upon, describing how individual ideas have typically converged to develop more complex ones. This section will conclude with the choice of *causal maps* as the specific type of solution architecture we will attempt to build with a crowd in this thesis. In Chapter 3, a series of pilots will be described, which enable a collection of design principles for a causal mapping platform to be suggested. Finally, Chapter 4 implements these design features in a prototype platform.



## 2 The Architecture of Solutions

### 2.1 Solutions Versus Solution Architectures

The goal of this thesis is to explore how members of a crowd can come together to build a *solution architecture* collaboratively, as opposed to offering individual solutions to a given question. To understand the advantage of such an effort, we must first begin by explaining the concept of solution architectures, which refer to how individual ideas are stitched together in order to develop a larger, more complex idea. These ideas can be stitched together in different ways, as a function of both what the ideas are, and how they depend on one another. The resulting structure of ideas is known as the solution architecture.

#### **Component vs. Architectural Solutions**

A noteworthy way in which solution architectures have been conceptualized in the literature is by Henderson & Clark (1990). There, in an attempt to explain the phenomenon of “radical innovation”, they differentiate between “component” and “architectural” knowledge. They look to the development of semiconductors in their examples, and highlight the difference between “component” knowledge – which refer to the individual components of an overall product, such as the blades or housing of a fan; and “architectural” knowledge – referring to the way in which these components are actually linked together. As such, different architectures can arise from the same components, where in some cases the blades of the fan may be inside the housing (such as in an air conditioning unit), and in some cases, outside (such as in a ceiling fan). An overall solution architecture can therefore be defined by the components it comprises, and the interdependencies between these components.

#### **Solution Architecture in System Architecture Literature**

The same principles apply in system architecture, as introduced by Crawley et al. (2016), which defines a complex system as one with “many elements of entities that are highly interrelated, interconnected, or interwoven”, and where these entities “are themselves likely to be systems”. Developing successful solutions to a given goal requires creating a concept of how different forms (i.e. the “physical or informational embodiment of a system”) should be interrelated so as to achieve specific functions (i.e. an “activity, operation or transformation that causes or contributes to performance”).

Crawley et al. recommend that in order to arrive at such solutions, it is important to build up options of “concept fragments”, which are constituent ideas within the overall solution, which link together several forms and the functions they are meant to achieve. These fragments then are linked together to form larger and larger fragments (by making all “possible combinations of the fragment”) until an overall solution architecture is achieved.

#### **Solution Architecture in Organizational Literature**

The conceptualization of solution architecture as ideas and their interdependencies has also been utilized widely in the organizational literature. March (1991) portrayed an organization’s knowledge as a binary string of  $m$  dimensions, which would produce different values depending on their constituent values. Levinthal (1997) then extended this to include the specific

interdependence of ideas, by allowing for the value of a single idea to either increase or decrease depending on what other ideas it was connected to.

In the context of online organization, the Collective Intelligence Genome (Malone et al., 2010) would classify this type of solution architecture as a “collaboration” – pertaining to crowd tasks which cannot be conducted entirely independent of one another.

### **Overall Description of Solution Architecture**

Overall, many interpretations have arisen for the concept of solution architecture. Fundamental to all, however, is the notion that solution architectures comprise both individual modules of information, and the interdependencies between them, which constitute a greater whole. Finding an effective solution architecture thus requires not just having the right pieces, but also knowing how they should come together.

## **2.2 Leveraging Open Innovation to Build Solution Architectures**

Current open innovation tends to be designed to solicit individual solutions from a crowd, as opposed to solution architectures.

### **Structured Solutions From A Crowd – Innocentive and Lego**

Innocentive, for example, is a platform where companies can pose questions to a crowd, such as the request to propose the “design of [an] affordable... off-grid lighting device” (“Dual Use Off Grid Illumination Device”, 2007). Constraints for such a design are provided along with the request, and members of the crowd are invited to contribute design solutions (Lakhani, 2008). Ultimately, the proposals are selected between by the organization hosting the competition, and winners are granted intellectual property and cash remuneration for their efforts. Individuals can choose to work on joint proposals; however each proposal does not build on others. Lego ran a similar type of open innovation effort, where consumers could innovate on new Lego designs, the best of which were then turned into actual products (Lakhani et al., 2012).

### **Freeform Solution Architectures From A Crowd – OpenIDEO**

A similar format is applied in OpenIDEO, where the crowd is leveraged to find solutions for social challenges, for which they are rewarded with reputation points as opposed to money (Lakhani et al., 2013). On this platform, unlike Lego or Innocentive, individuals or teams are able to submit a wide variety of content, such as ideas, photos or online resources they feel are relevant to the challenge posed. They are also able to solicit feedback from other users who provide comments. As such, far more curation occurs than on Lego or Innocentive, where OpenIDEO moderators must assist in synthesizing submitted ideas, and reframing promising solutions to the crowd to invite further ideation.

### **Structured Solution Architectures From A Crowd – TopCoder**

Another notable example is TopCoder, where a programming task is given to the crowd (Lakhani et al., 2012). Initially, the platform leveraged the crowd to not only provide solutions, but to first conceptualize given problems better, then specify the unique subtasks within them which were used to solicit solutions. Over time, however, the costs of conceptualizing and specifying tasks were considered too high for companies using the platform, and it has shifted more towards a solely solution seeking (as opposed to question refining) platform through the

introduction of TopCoder Direct, making it similar to other typical open innovation efforts (Lakhani et al., 2012).

### **Platform Difficulties**

These few cases give examples of how open innovation has been a useful way of soliciting the best ideas from a crowd, based on their answers to a given problem. It has been more challenging, however, to have ideas that interact with one another – be it freeform ideas that grow through discussion and require heavy synthesis on the part of the platform, or ideas which modify the ideas which come after it (such as TopCoder’s conceptualizations and specifications affecting the type of subtasks given to other users).

### **Why We Need Better Solution Architecture Platforms**

Thus the challenge still remains to develop an open innovation platform that can assemble solution architectures at scale, and in a value-adding manner. One can foresee that such a platform might prove beneficial in assembling the varied knowledge of a crowd into rich and coherent solution architectures; building upon the diversity of ideas as opposed to having them compete against one another, and better preserving participants’ efforts as a result. With problems as complex as they are these days, it is unlikely to expect one individual to come up with all of the right component knowledge and their requisite linkages in a solution architecture; instead, the component and interdependence knowledge of multiple participants should ideally be leveraged in generating solution architectures. One can also envision that with a crowd able to develop more intricate solution architectures, companies can pose even broader questions to the crowd. Alternatively, crowds may even be able to better self organize around projects they care about, rather than relying on questions being posed by a larger organization.

Overall, such a platform would thus foreseeably produce higher quality solution architectures, at a lower cost to the effort of the crowd. To elicit such advantages however, we must first understand how it is that solution architecture is in fact effectively assembled, and how to achieve it at a minimum level of effort from either the crowd or platform moderators.

## **2.3 How Solution Architectures Are Integrated**

In order to design a platform where crowds can build solution architecture together, it is necessary to first understand how such integration occurs. How is it that contributions from a crowd can be solicited so as to be interdependent on one another, and grow into a cohesive architecture, as opposed to disjointed solutions to a problem?

### **Putting Many Ideas in the Same Place**

At one extreme, no integration occurs – individuals might randomly submit their ideas to a question such as “How can carbon emissions be reduced”, and produce answers such as “increase consciousness about climate change” or “plant more trees”. Integration might then be done by reviewers or experts after the fact, who consider how all of these ideas might be implemented together. However, because these solutions can be of any shape and size, the answers might not all complement one another, and it may not be possible to conduct them all simultaneously (e.g. increasing consciousness about climate change might warrant making the climate situation more dire, which may just involve razing down more trees in the short term).

This is similar to the case of OpenIDEO, where much moderation is required to integrate contributions into useful solutions.

### **Integrating Categories of Ideas**

One way in which these disjointed solutions might perhaps be more easily integrated, and allow for more complex solutions to result, is by requiring that the answers given fall within certain categories. As such one might ask that the crowd either contribute solutions of the “technology”, or “policy” kind, making it easier for an external integrator to group complementary solutions together into an architecture. If “technology” were even further separated into “manufacturing technology” and “consumer technology”, an integrator might be able to pull contributions from the “manufacturing”, “consumer” and “policy” categories that are more likely to be cohesive, as opposed to three likewise random draws from the earlier format of an open bucket of solutions.

It is for this reason that it can be very helpful to leverage pre-existing models for solution architecture. One such example is Object Process Modeling, which separates potential ideas into categories such as “agents” and “processes”, to distinguish a group of “agents” that are able to bring about “processes” such as “transforming” or “consuming” a resource. The Climate CoLab is an example of another platform that solicits climate change solutions from different categories such as geographies or applications (e.g. groups such as “energy” or “buildings”), and allows subsequent users to aggregate these solutions into a larger solution “family”, so as to achieve better solutions to complex problems. (Malone et al., 2017)

What all these efforts represent is essentially the value of problem reduction, and the importance of asking the right question. If one is able to reduce the problem of lowered carbon emissions into the problems of “limiting gas usage”, “building up the battery industry”, and “changing mindsets towards climate change”, then specific questions can be asked in each of these categories, which would allow their solutions to then be integrated into a cohesive solution architecture.

### **Integrating Without Having to Presuppose Architecture**

The biggest drawback to this, however, is that it *requires the ability to reduce the problem*. In other words, it requires that the organization posing these questions in fact knows enough about the problem at hand to be able to reduce it to a set of sub-problems – that it *already knows the solution architecture*, and is merely missing the right components to fill it with. By reducing the problem in this way, the organization effectively reduces much of the benefit that the crowd can provide, negating their ability to generate entirely novel solution architectures which the organization is unable to think of.

There must therefore be a way that the crowd can somehow come to build an architecture of its own, and be able to interface itself with other ideas in the crowd so as to build upon them. It is ideal, for example, that if one suggests “planting more trees”, another can contribute an idea which builds on that by suggesting “creating more natural reserve lands”, and yet another can suggest “giving companies tax breaks for planting trees”. All of these ideas would thus build upon one another, without the necessary constraint of a solution architecture template.

### **Relational Value of Diverse Ideas**

How is it then, that ideas can build upon one another? Given one idea, how can we invite other members to build on it ways that leverage their creativity and diversity of perspectives, to build an architecture of solutions that no one person could have developed alone? How can we leverage both the crowd's diverse knowledge of components, and how they can be linked?

The literature on diversity has some insight as to how the perspectives of different individuals can provide the most relational value to one another. Dahlin et al. (2005) saw diversity as a route to increased learning, by touting the value of the “depth” of information, where team members would “trigger each others’ knowledge”, enabling them to learn more. They do not go into the details of how this is possible, but it is perhaps similar to Cohen & Levinthal (1990), who espouse diversity as a means of creating greater “absorptive capacity” for new knowledge. This was perhaps also similar to Amason’s (1996) recommendation of the value of diversity for “decision quality”, where increased diversity of perspectives would force consideration of underlying assumptions, and produce higher quality decisions.

As such, because these mechanisms are unclear, the optimal way in which to enable the integration of solution architecture is unknown. The subsequent chapter summarizes a series of pilot experiments conducted to discover how in fact this might be achieved in an online platform.

## **2.4 Types of Solution Architectures**

There are many different kinds of solution architectures that might be created – comprising different kinds of individual modules, and different ways in which each of these might be related. System Dynamics, for example, is a type of architecture in which modules are “stocks” (of materials, knowledge, resources, etc.) while the relationships between them are “flows”, which can be positive or negative. Another type of architecture is the Object Process Model (OPM), where modules can be broken down into categories such as “objects”, “agents”, or “instruments”, which are related to one another by “processes”, which comprise actions such as “transformation” or “consumption”.

This thesis, however, will focus specifically on the architecture of causal maps. This is because it is hypothesized that causality may be a mechanism capable of defining a consistent interface by which ideas might integrate with one another, yet being generic enough so as to not impose any type of solution architecture onto the integration of pieces.

The type of causal relationship in mind is the most generic one – where two events can be related in an architecture simply because one is a cause of another. Many types of modeling in fact fall within the concept causal maps, but are termed differently because of how the types of causes have been categorized. In OPM for example, while agents and instruments are both treated differently, they are both essentially causes. The way in which they cause their effects have been divided into “transformation” and “consumption”, even though both of these fall within the overall category of cause and effect relationships. Other ways that one can divide types of causes is into “why, who, what, where and how”; as sufficient or contributory causes; or “material, formal, efficient, or final” causes (Aristotle).

For the sake of this thesis, however, we are interested in the most generic type of causality, with the hope of creating a map that links ideas into an overall architecture, only by virtue of the fact

that they might be causes or effects of one another. Such generic causal maps have been used extensively before – often recommended as means of understanding phenomena when no theory exists yet, or assisting in decision making by understanding the possible outcomes of a decision (Narayanan & Armstrong, 2005).

In this thesis, the aim is to explore the particular application of causal mapping as a means of integrating novel solutions for climate change. Climate change is an extremely complex problem, touching upon many facets of society, geographies, industries, and ways of life. As such any potential intervention to mitigate it will have to consider many elements, far beyond what a single individual may be able to consider on their own. Tackling climate change is therefore a quintessential scenario in which solution architecting might prove useful.

## 2.5 Platform Goals

The ultimate goal is to work towards creating a useable causal mapping platform which will enable members of the public to come together to address climate change issues. The benefit of such a platform will be evaluated along the dimensions of **usability** and **usefulness** – where usability is the extent to which new and inexperienced users might be able to access the platform and use it to contribute solutions to an overall architecture; while usefulness refers to the quality of the overall solution architecture generated by the crowd. Usefulness will specifically be operationalized as the **creativity** and **impact** of the solution architectures generated.

It is recognized that these measures of usability and usefulness can in fact be quite subjective, and in fact they will remain as subjective measures for this thesis. This thesis is very much an exploratory one. Given the lack of prior research in both the factors material to online collaboration in a causal context, or how such a platform for collaboration might be implemented, there are too many unknowns to warrant an objective exploitation of the subject at hand. As such many of the ideas in this thesis have to be developed in an exploratory manner (e.g. using pilots instead of experimental studies, or the development of prototypes without a strong rationale for the features they include).

### 3 Discovering the Dynamics of Collective Causality

As stated in the prior chapters, in order to build a successful platform for crowds to assemble solution architectures, it is necessary to understand how such architectures become integrated from the discrete ideas offered by individuals. The literature is unclear, however, as to how exactly this happens.

A series of pilots are conducted to understand what variables are material to the ability of a crowd’s ideas to integrate. Specifically, these pilots are conducted around the task of collaboratively developing a causal map to address climate change issues. At each stage, a set of takeaways will be reasoned from the results of the pilots, so as to inform potential design features which a collaborative causal mapping platform might ultimately include.

#### 3.1 In-Person Collaboration

Most of the studies in this chapter are done using Amazon Mechanical Turk, such that testing is done amongst strangers online – the eventual target context for this platform. However, an initial study was done with an in-person group, so as to better understand how the mechanics of solution architecting might work in a real world brainstorming discussion, to determine what added benefits of this environment might be adopted online as well.

Table 1. Summary of Pilot #1 parameters

<b>Pilot #1 – In-person Collaboration</b>	
<b>Number of tasks</b>	1
<b>Number of contributors per task</b>	4
<b>Amount paid per task</b>	0

A set of guidelines were provided to the team for their discussion, where they were asked to brainstorm a set of enabling actions, as well as the individuals or groups who might be responsible for making them happen. The effect of these guidelines was to provide a basic architecture for the solutioning process, as such the situation was not a freeform real world, in-person brainstorming environment. The aim, however, was to implement an architecture which might be similar to the causal architecture eventually desired in an online platform (as described in the prior chapter), and thus to see how the in-person setting would influence such a process. The goal was for the pilot to identify important features which a causal solution architecting process would have to incorporate from an in-person context.

#### **Topic Addressed**

This pilot was conducted with a group of graduate students undertaking a capstone project to implement a clean energy system in Ecuador. The team was brainstorming solutions to the following question: “*How can we improve access to clean water and energy in Ecuador?*”

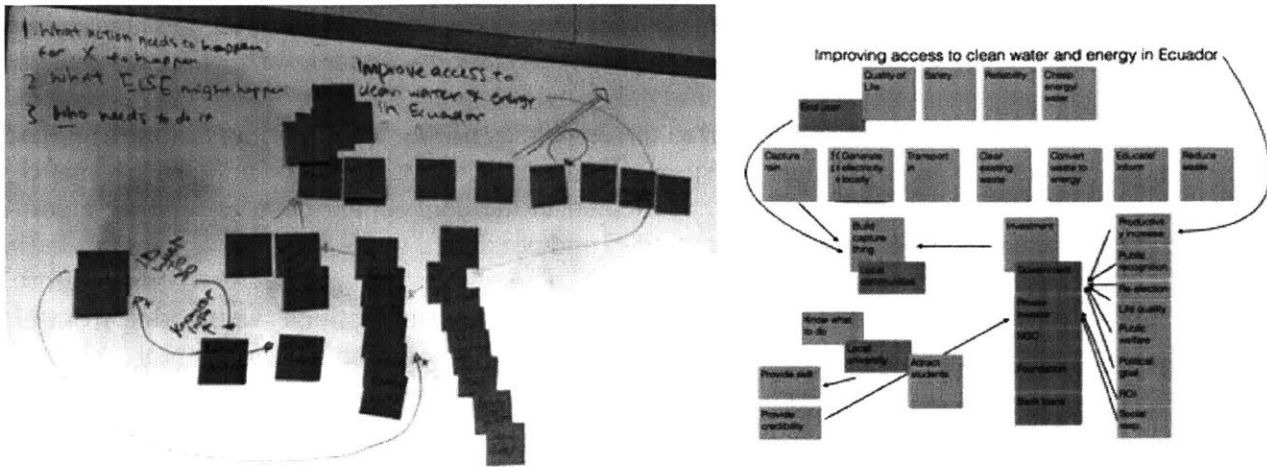


Figure 1. Eventual causal map output

### Solutioning Guidelines

In order to mimic as much as possible the modular form in which ideas would be contributed in an online context, members of the team were asked to contribute their ideas using post-its. As such, the format of the group brainstorm was where individuals added to, and removed, post-its from a central board. As they did so, they were asked to consider (1) What action needs to happen for a given action (i.e. an existing post-it on the board) to happen? (blue post-it) (2) What other actions might need to take place as well? (blue post-it) (3) Who will need to perform such an action? (red post-it)

### Team Roles

The team naturally split itself into four roles, where (1) one was the team leader and the primary contributor of post-its; (2) another was responsible for adding these post-its to the board; (3) a third was a more occasional post-it contributor, and (4) the last documented findings from the process.




Figure 2. Group undergoing brainstorming discussion to develop causal map



### 3.1.1 Analysis

A basic coding exercise was done to categorize the responses of the team members into different types of key interactions (e.g. an idea, versus a clarification about an idea), and the extent to which ideas raised by the group were in fact reflected in the contributed post-its (e.g. a person may have mentioned an idea, but it was not captured and not added to the board).

Over the course of a 20 minute discussion, 63 question-relevant interactions were coded from a video recording, which included 50 recommendations – comprising both new ideas (e.g. “safety would also be something that the government might want”, or “build a pipeline, that would be one way of transferring”) or the recommended removal of ideas, and 13 clarifications (e.g. “philanthropists? Or is that like NGOs”, or “wasn’t that what you meant by...”). Of the 50 ideas, 41 were recorded on the board.

The ideas were also coded according to the level of the causal map at which they were added – referring to whether they were added proximally to the overarching question of how to improve access to clean water (i.e. a “level 0” idea), or added as a cause of a given cause (i.e. a “level 1” idea), and so on and so forth. The team discussion was orchestrated such that they were free to add ideas at any level of the map created. The discussion initially began with causes being added at increasing depth (i.e. adding a series of level 1 ideas, followed by a series of level 2 ideas), however this pattern did not persist. Of the 41 ideas recorded, 11 were added a more proximate level than the idea preceding it (e.g. a level 1 idea added after a level 4 idea).

An important observation was also that over time, the team began to move away from the given guidelines of how to link solutions to one another, and in fact added their own architecture of how ideas should interact. This is likely because the team was well versed in developing Stakeholder Value Networks (a manner of identifying stakeholders and the types of value which they provide one another), as well as the methodology of Object Process Modeling (wherein ideas are divided into different types of processes or transformations; actors are divided into human agents of instrumental objects). Over the course of the discussion, the team reverted to the types of terminologies established in these established solutioning processes (e.g. making clarifications such as “but rain is an operand”).

As such, given the familiarity of the team with solution architecting, the subsequent takeaways may not be applicable to groups in general, who are more likely to have difficulty integrating their thoughts with one another into a coherent overall architecture. (The amount of guidance that might be required in such a scenario is explored in the next series of pilots.)

### 3.1.2 Takeaways

The exercise was essential in identifying some key dynamics the team used as they worked towards a solution.

The first was that there is a cost to writing down every idea, such that it was sufficient for 9 of the 50 ideas to be discussed and refuted, before they were committed to a post-it. This indicates that some important ideation may occur which is too rapid to be captured in an online format, where ideas go through a slower frequency of being declared and exchanged with one another.

Some studies, however, have indicated that increased reflection by each individual before submission of an idea results in greater creativity of the ideas submitted (Taylor et al., 1958, Diehl and Streobe, 1987).

The level of clarifications in the discussion demonstrated the difficulties of integrating individual contributions so as to result in a cohesive solution architecture. However, as explained above, the team was already familiar with other solution architecting methodologies. In addition, they also had a great degree of familiarity with the topic at hand, having worked on it for several months already. As such, it was apparent that the team had a strong tacit representation of the model which they shared, implying that an online group with no such prior experience would be less likely to generate as coherent a solution architecture. Indeed, having a common mental model of a situation has been found to improve team’s creativity and innovation (Drazin, Glynn, & Kazanjian, 1999; Dutton & Dukerich, 1991).

It can be expected that the same exercise conducted with strangers online might involve many more erroneous contributions (e.g. ideas which are not actually enabling actions) and the need for clarifications. Hence this pilot indicates the necessity of an equivalent online platform to have either clearer instructions as to how contributions should be made, potentially by educating users on a modeling format with some properties of OPM. Perhaps another way of achieving clarity would be through a separate platform where individuals could discuss the meaning of contributions; or perhaps a means through which to indicate the validity of others’ contributions.

Lastly, given the extent to which ideas were added at varying levels of the causal map, a key functionality would be to enable additions to any part of the overall map, as opposed to only its lowest levels. This will allow the entire architecture of the solution to be constantly changed, as opposed to only changing the latest additions to it. Perhaps users might also retain visibility of the whole map, in order to select where they can make more valuable contributions. The parts of a problem that users attend to, however, can have significant impact on the creativity of the output (Ansburg & Hill, 2003).

The table below summarizes the takeaway from this pilot, comprising the set of features which may be considered for addition to the eventual platform.

Table 2. Takeaways from in-person pilot

Proposed Features	Potential Pros	Potential Cons
<b>Personal refinement:</b> Ability to refine ideas separately	Higher quality ideas	Lower volume of ideas
<b>Leveraging existing methodology:</b> Educating users in a common modeling methodology (e.g. OPM)	Reduce confusion, improve creativity	Higher barrier to entry to contributing
<b>Discussion platform:</b> Some means by which users can clarify or validate others’ contributions	Increases quality of ideas	Increases complexity of platform
<b>Total map access:</b> Ability to add to any part of the causal map	Allows any part of the solution to be edited	Constrains contributions by anchoring ideas against a very rich solution

### 3.2 Online Collaboration – Freeform versus Guided Map

The in-person pilot indicated that much guidance is necessary to facilitate a smooth exchange of ideas. The subsequent set of studies thus experiment with the extent and type of guidance which may be required in an online exercise, in order to provide users a similar ability to build upon one another's ideas.

#### 3.2.1 Freeform Collaboration

The first online study examined a completely unguided context, to confirm the hypothesis that individuals would be unable to accumulate their ideas without any guidance at all. The prior study found that individuals were able to share their ideas in person, but largely because of the guidance provided by the researcher on what each contribution should entail, as well as their own prior familiarity with solution modeling techniques.

Although the goal is to work towards an online causal map for collaboration, this pilot is conducted in a *non-map* setting. It was hoped that understanding how ideas accumulate in such a setting would help to accentuate the benefits of using a map to integrate individuals' ideas.

Table 3. Summary of Pilot #2 parameters

Pilot #2 – Online freeform collaboration	
Number of tasks	3
Number of contributors per task	3
Amount paid per task	\$1

The study directed Mechanical Turkers to a Google document, where they were asked to develop a central proposal together. The document began blank, allowing the first user to contribute ideas from *carte blanche*. It was hoped that subsequent users would then develop a proposal and build upon one another's work as more users came to contribute. The prompt provided to users on Mechanical Turk was as shown in the figure below.

In September 2016, California signed legislation to reduce its greenhouse gas emissions to 40% below 1990 levels, before the year 2030. At present, the state is on track to meet its 2006 target to meet 1990 levels by 2020. Read more about California's new legislation and its efforts [here](#).

Use the area in the editable document (included in this folder) to brainstorm about what series of actions (such as policies, technologies, community initiatives, etc.) will enable California to reach this goal. What direct actions or indirect actions can cause this goal to be achieved?

**Instructions:**

- Write your answer in prose form
- If there already is an answer, edit and add to what is there, but maintain the prose form.

Figure 3. Prompt for users to contribute to a communal proposal

The following figures give two examples of responses obtained. In Figure 4, it can be seen that despite the given instructions, users added another proposal altogether, without modifying the

previous entry. This might have been a function of users being unclear about the given instructions. However, in the response of Figure 5, where it appears that subsequent users added to a central proposal, the editing history in the Google document indicated that they in fact added sequential lines to the given proposal. There was also little continuity between the ideas presented, as different contributors made no effort to logically link their ideas to prior ones. As such it was akin to Figure 4's list of ideas.

**Proposal:**

**Award companies that create innovative technologies which use less emissions. For example, give out yearly monetary awards to the company that has the best turnaround in reducing their emissions. Another example would be to give government funding to research into more global warming friendly innovations that have less carbon emissions. Avoid punishing people for emissions or raising taxes. Make it reward-based and keep things positive to inspire people and businesses to participate.**

**Proposal:**

**Improve mass transit from the outlying areas of LA to the city. LA is a huge metropolitan area with multiple edge cities as well as large suburban areas far away from downtown (i.e. The Inland Empire). This layout results in a lot of commuting traffic. Improving mass transit would take cars and the resultant greenhouse gases out of the environment.**

**Proposal:**

**Provide more mass transportation, use electric means to cut carbon transmissions from cars. Institute a tracking means to show which areas are using more electricity. Give tax credits for installations of smart appliances. Employ more solar panels in businesses and homes and provide incentives for using them.**

Figure 4. Sample output #1 of freeform collaboration, where ideas are listed separately

**Proposal:**

**I think that we should institute laws that require people to recycle. There are already government employees that go around making sure we put our trash cans out on the right curb, why shouldn't they make sure that we are doing the right thing by recycling as well? We could all be assigned different containers for glass, paper, etc. And just do the right thing by sorting everything, making more of our products out of recycled materials will do wonders for our environment and for sustaining our resources, especially those which aren't very renewable.**  
**I also think we should institute reducing the number of cars using fossil fuels. We should set a target date, say 10 or 20 years, and mandate that by that date all vehicles on public roads must be electric or otherwise using renewable forms of fuel. By reducing the number of smog emitting vehicles, we can greatly reduce greenhouse gas emissions permanently and start to bring our environment back to its former greatness.**  
**Additionally, we should create tax credits to encourage the use of carbon offsets or carbon credits. These offsets**

Figure 5. Sample output #2 of freeform collaboration, where ideas are in the same prose, but ideas remain independent of one another

### 3.2.2 Freeform Alternative

Presuming that the results obtained were possible because of a poor prompt (where users were asked to “edit and add”, and not necessarily create a cohesive overall argument), a follow up

pilot was conducted, where users were asked to ensure that “the whole output should make a cohesive argument”.

Table 4. Summary of Pilot #3 parameters

Pilot #3 – Alternative freeform collaboration	
Number of tasks	1
Number of contributors per task	5
Amount paid per task	\$0.03

After each user’s contribution, the prose they improved upon was offered to the crowd again for the next user’s improvement. Figure 6 shows the prompt given to the user at one of these stages, where they are asked to improve on an answer.

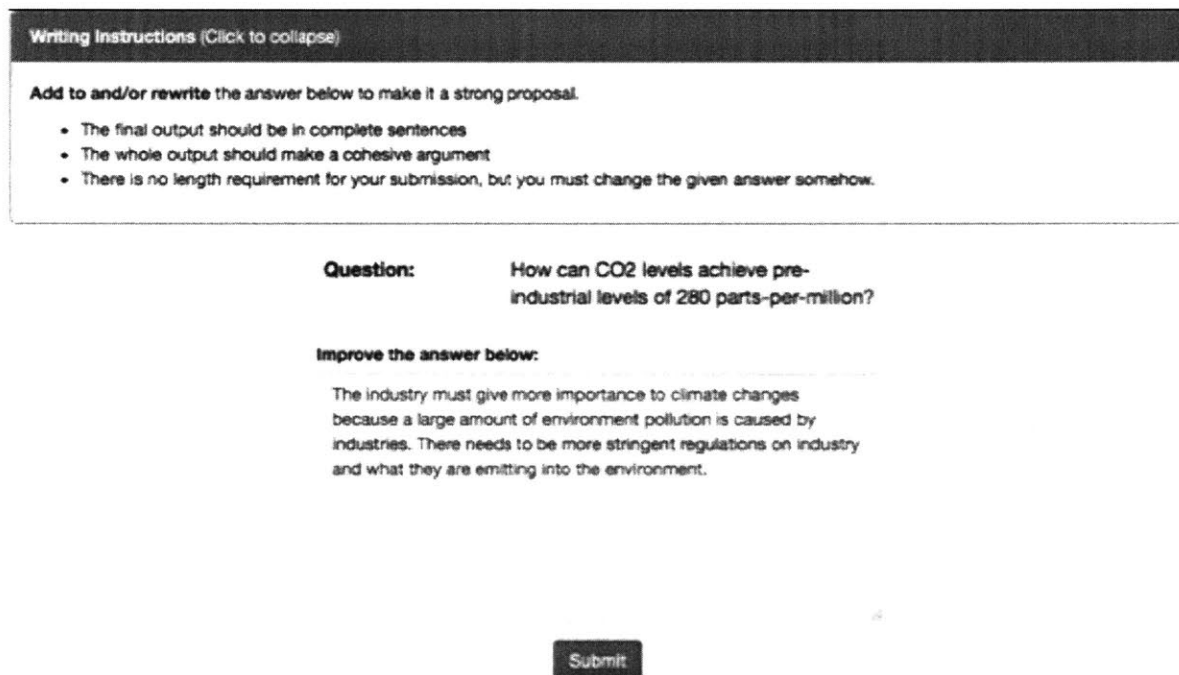


Figure 6. Example prompt to improve on a given prose

### User contributions

Using this format, each user’s improvements were indeed able to amount to a cohesive whole. As shown in Figure 7, two improvements (v2 and v4) can be considered to have added new concepts (i.e. why industry would care more, and what actions could buttress this). The other types of modifications were linguistic, with v5 improving the clarity of v4’s statement, and v4 editing out v3’s contribution which was difficult to understand.

### Linguistic Versus Conceptual Cohesion

Ultimately, the goal is to create a proposal that is conceptually cohesive, as opposed to linguistically cohesive – where the emphasis is on creating a solution architecture where individual ideas add to the value of one another. The added linguistic contributions are seen to be a feature of asking for a “cohesive argument”, and it is supposed that by framing the guidelines better, and using a map format instead of a freeform one, contributions will likely be more

conceptual than linguistic. As such this reinforces the choice of integrating using a map as opposed to communal prose, so that we can distill the crowd’s effort into conceptual elements of a solution architecture.

	Output
v1	Industry start taking climate change seriously
v2	The industry must give more importance to climate changes because a large amount of environment pollution is caused by industries.
v3	The industry must give more importance to climate changes because a large amount of environment pollution is caused by industries. yes environment does not pollute then climates all right
v4	The industry must give more importance to climate changes because a large amount of environment pollution is caused by industries. There needs to be more stringent regulations on industry and what they are emitting into the environment.
v5	Industries must realize the immense importance to climate change, because a vast amount of environmental pollution is caused by their emissions. In addition, these emissions need to be regulated more stringently by the federal government.

Figure 7. Subsequent improvements upon the proposal by each user.

All the same, linguistic cohesion can be useful in asking users to make sense of prior contributions. For example if they are written in different tenses, or simply have incomprehensible language, this can obscure the potential conceptual value of the contribution. As such, features to improve the language in contributions would be useful. One way that this can be done is by having users fill in the blanks of a sentence, such as “a cause for the above event is \_\_\_”, so that contributions are at least structured in a similar manner. Another way this might be achieved is to seed the map with sample nodes, so that users have an idea of the type of language that is expected.

### 3.2.3 Guided Map Collaboration

Having understood the opportunities and constraints inherent in a freeform online collaboration, the next pilot explored how the same question from Pilot #2 (the online freeform collaboration) might be answered differently in a map format.

Table 5. Summary of Pilot #4 parameters

Pilot #4 – Online Map Collaboration	
Number of tasks	1
Number of contributors per task	3
Amount paid per task	\$1

A website was developed using Meteor (a javascript-based web development platform), where users could come to contribute ideas to a map. Users were able to contribute nodes, edit and delete them, as well as create linkages between them. This would enable users to interact in a similar manner to what in-person collaborators would do with post-its on a board. Mechanical Turkers were directed to this map to add their contributions.

### Guidelines Provided

Specific guidelines to structure action on the site were also included – comprising the constraint that individual cause contributions should be between 50 and 200 characters long, and that users must indicate whether causes had a sufficient or contributory causal relationship with the target node. Users could choose to add their own node, or create novel linkages between existing nodes. The prompt provided to the users, as well as the basic features of the website, are shown below.



In September 2016, California signed legislation to reduce its greenhouse gas emissions to 40% below 1990 levels, before the year 2030. At present, the state is on track to meet its 2005 target to meet 1990 levels by 2020. Read more about California's new legislation and its efforts [here](#).

Use the map below to brainstorm about what series of actions (such as policies, technologies, community initiatives, etc.) will enable California to reach this goal. What direct actions or indirect actions can cause this goal to be achieved?

**Note:** If you do not see a map below, refresh the page.

---

**Propose a cause for** **Test Node A. Sample causality** **would be entered here):**

**New cause:**  **Existing cause:**  
 E.g. technology, policy, i or **Click on a node**

**This cause is:**  
 **Sufficient (it alone is enough to cause the outcome)**  
 **Contributory (it helps the outcome to occur)**

**Edit a cause**  
 Click a node's **edit** **delete** to edit/delete its content

**History**  
 • Fri Apr 28 2017 17:33:53 GMT-0400 (EDT): dincopasta proposed that "Test Node A. Sample causality would be entered here" can cause "By 2030, California reduces its greenhouse gas emissions to 40% below 1990 levels"

Figure 8. Basic causal mapping platform developed for the pilot

## User Contributions

Three people contributed to the map, as was the case in the initial freeform pilot. All users were able to effectively connect the ideas causally, in that each node was in fact a foreseeable cause for the cause it pointed to. The users' contributions to the eventual map are color coded in Figure 9. The first contributor provided the bulk of the causes, and used the platform to constantly edit some of the ideas he created. The next contributor added a cause and a causal link between two existing causes, while the last edited the words within the first user's cause.

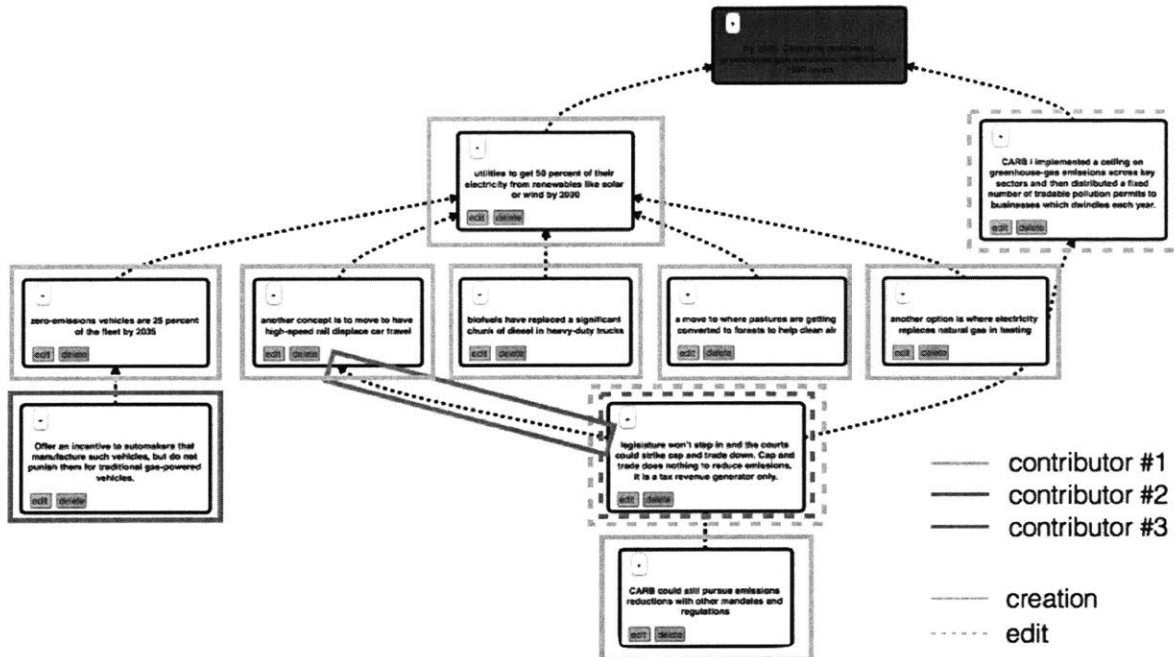


Figure 9. Users' contributions to a collaborative causal map

### Feedback Provision

This last edit is noteworthy in that the user added a sentence to the prior node, to presumably explain why the prior contribution might be irrelevant, or need further refinement. The initial user had contributed the cause “legislature won’t step in and the courts could strike cap and trade down”. The subsequent user added to this that “cap and trade does nothing to reduce emissions, it is a tax revenue generator only” – which appears to be a commentary that the initial user’s contribution was irrelevant. This user perhaps interacted this way because while there are ‘edit’ and ‘delete’ features on the website, there is no mechanism to provide feedback to another user, to enable the latter to improve his contribution. This pilot thus suggests the possibility of such a feature for future consideration.

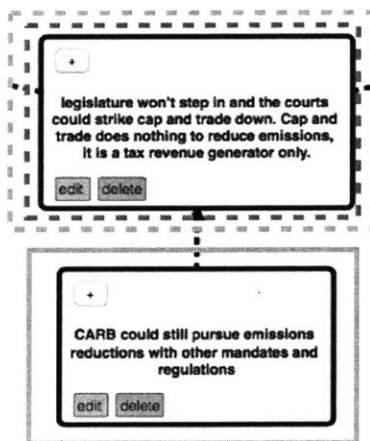


Figure 10. Lower node cannot be understood independent of upper node



### Standalone Comprehensibility

It should also be added that some contributions were difficult to understand independently, such as the contribution that “CARB could still pursue emissions reductions with mandates and other regulations”, shown in Figure 10 above. It uses the terms “still” and “other”, which indicate that it is referring to another contribution, and indeed it can only be made sense of when taken in context of the node it points to. This shows a lack of a clean causal connection between nodes, since this particular relationship is perhaps of two events which can occur in place of one another, rather than one causing the other.

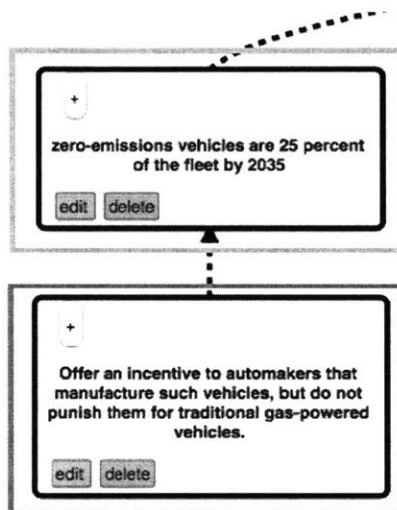


Figure 11. Node below actually includes two integrated concepts. In addition, it cannot be understood independent of node above.

Likewise, another user contributed the idea of “offer an incentive to automakers that manufacture such vehicles, but do not punish them for traditional gas-powered vehicles” as shown in Figure 11. Similar to before, this response requires an understanding of the node it points to, in order to understand what is meant by “such vehicles”.

### Complexity of contribution

The lower node in Figure 11 also raised a fairly complex contribution which actually comprised two ideas – that there should be “an incentive to automakers”, but that the value of this is perhaps contingent on there not being a “[punishment]... for traditional gas-powered vehicles”.

Such a contribution is not anathematic to the concept of a causal map. In fact, it is reasonable to expect that any architecture of solutions is going to include individual solutions that are themselves architectures of smaller ideas. Indeed, the idea it points to, “zero-emissions vehicles are 25 percent of the fleet by 2035”, itself integrates ideas of type of technology, proportion, and timing.

These contributions raise two potential shortcomings of the current map format. Firstly, they indicate that there may be crucial relationships between ideas that we have neglected, such as the conditions under which a causality may be true or not. Allowing for such complex relationships between contributions would indeed allow more complex solution architectures to be assembled, without necessarily posing an overly constraining solution architecture on the whole, as was the

fear with overly proscribed platforms in Chapter 2. These relationships of complementarity or conditioning can be expected to be as problem-agnostic as causality is.

A second shortcoming is that the more complex an idea, that longer one might take to understand it, as well as propose a cause for it. Potentially useful ideas may therefore become dead ends in the causal map, simply because users would prefer to contribute to more simplistic alternatives. Some strategy might therefore be employed to force the simplicity of the idea involved.

The extent to which these two issues may affect the quality of the overall solution architecture is unclear – we are forced to weigh the benefit of richer parts with richer connections constituting a richer whole, versus the possibility that such complexity may limit the integration of a whole. This challenge will not be tackled in this thesis. Instead, it will be assumed that because this platform would be introducing a novel format of crowd sourcing, that the lower the barrier to entry to use the better. As such, the simplicity of each node is desired, and a mere causal relationship will be used.

Ways in which the simplicity of the nodes might be encouraged is to reduce the character upper limit for contributions, such that individuals cannot enter overly complex ideas. Perhaps the lower limit (50 characters) was in fact too high, forcing users to overly embellish their contributions. In addition, as above, pre-populated nodes might be useful as examples to guide users on the type of concepts expected.

It should be noted that lower character limits can be helpful to eliminate nonsensical contributions. The map below shows contributions such as “BBBB” and “ZZZZZ” made by Turk users on another pilot with a similar goal, where no character limits were enforced.

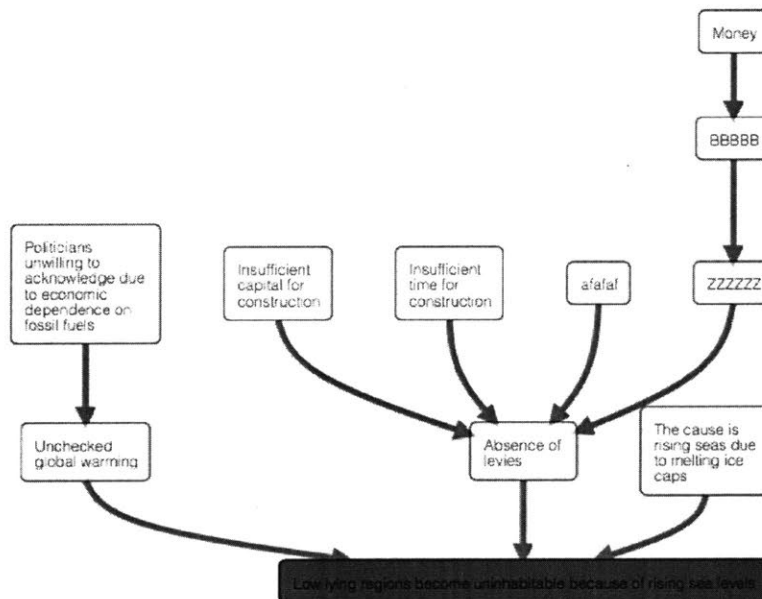


Figure 12. Crowd-generated causal map without minimum character limits

### 3.2.4 Takeaways

The online pilots help to clarify an additional set of potential features that can be added to the eventual platform, so as to distill the desired type of contributions (such as by enabling feedback on prior contributions) and guard against contributions which may either fail to build upon prior ones (as in the addition of disconnected ideas to a given prose), build upon them in trivial ways (as in the linguistic edits made) or be altogether incomprehensible.

The following table summarizes the potential features this exercise has suggested, and their associated pros and cons.

Table 6. Takeaways from online pilots

Proposed Features	Pros	Cons
<b>Feedback mechanism:</b> A way that users can enable prior users to improve their contributions	Improves the quality of each contribution	Can complicate the user's interaction with the platform
<b>Fill in the blanks:</b> Encourage linguistic similarity of nodes	Improves comprehensibility of each contribution	-
<b>Seed nodes:</b> Pre-populate the map with sample ideas	Encourages simplicity of nodes and commonality of language	Can anchor thinking and reduce diversity of contributions
<b>Control character limit:</b> Adjust the lower and upper character limits	Promotes simplicity, if bracketed correctly	Potentially excludes valuable ideas with too low of an upper limit, or makes it easier to contribute nonsensical ideas with too low of a lower limit.

### 3.3 Asking the Right Question

A common theme that appeared in the pilots above concerned the definition of the interface between contributions. While causality is used to connect ideas, the way in which this causal contribution is solicited is crucial to the success of the platform.

For example, the offline experiments suggested the importance of clear guidelines that might define what is meant by causality, and the varying effects expected of interacting with all prior contributions as opposed to just a few of them; while online experiments then suggested the possibility of using fill-in-the-blank prompts to improve the language of contributions.

As such, a subsequent set of pilots was devoted to the way in which contributions might be solicited – specifically, (1) how the prompt for a cause was written; (2) whether the presence of other contributions might anchor a user; and (3) whether users will have an availability bias in searching for aspects of a causal map to contribute to.

#### 3.3.1 Question Framing

Mechanical Turkers were used to test the relative effectiveness of three main styles of questioning.

Table 7. Summary of Pilot #5a parameters

Pilot #5a – Question Framing	
Styles of questions	3
Total questions across all styles	20
Number of responses per question	5
Amount paid per response	\$0.05

These styles evaluated were:

- **Simple cause:** Turkers were asked to contribute causes for a given effect. Within this style, different types of language were used to solicit this cause, such as “propose a cause for:” or “what is a scenario that can lead to the following?”. Within this, the target effect, for which the cause was to be found, was varied as well. Specifically, individuals were asked to propose a cause for a state (e.g. CO2 levels are at XX levels), or a process (e.g. CO2 levels are *reduced to* XX levels). The point of varying the target effect was to see if the quality of the contribution might somehow be dependent on it, as opposed to the way the cause was asked for.

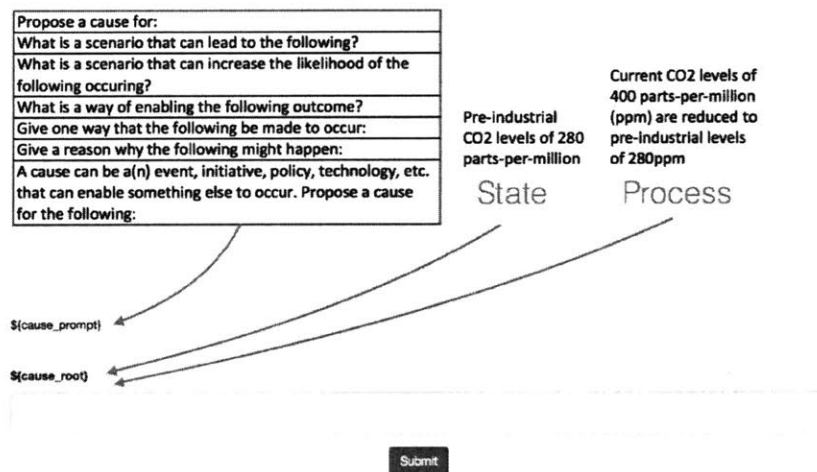


Figure 13. Prompts given to Turkers under "simple cause" condition

- **OPM:** The participants in the in-person pilot had shown much success in integrating their ideas. One suggested reason for this was their familiarity with Object Process Modeling (OPM), a highly structured format for system architecture. As such, this exercise was used to explore the extent to which users could be taught a light version of OPM, and the extent to which they could use it to suggest causes.

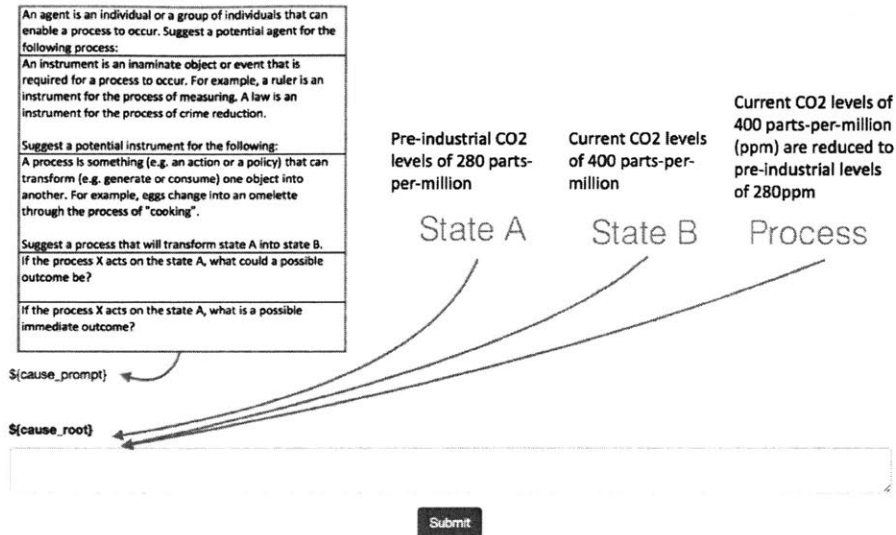


Figure 14. Prompts given to Turkers under OPM condition

- **Outcome:** Here, the causal relationship was reversed, where Turkers were asked to instead provide an *outcome* for a given state or process. As with the simple cause style, the language used for this question was varied, such as “What is a potential outcome for the following?” or “If the following occurs, what is likely to happen next?”

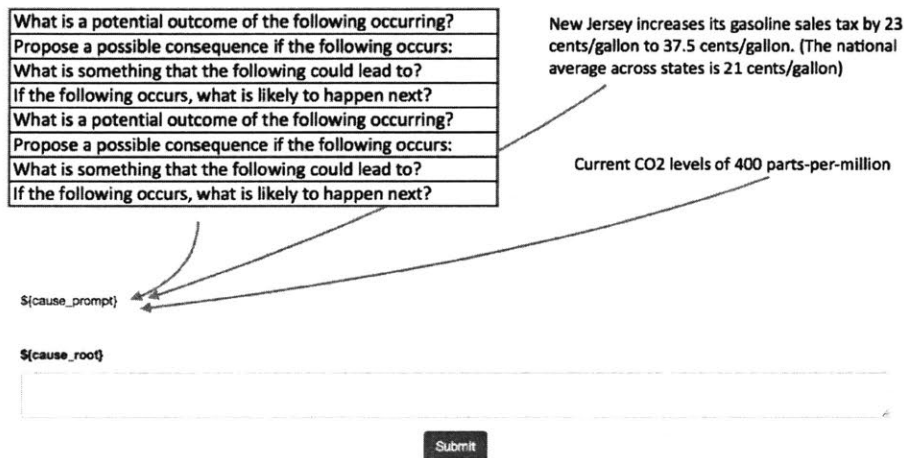


Figure 15. Prompts given to Turkers under "Outcome" condition

### Outcome measure

In order to evaluate the relative benefit of the different question framings, a separate Turk task was used to evaluate the relatedness between the responses from Pilot #5a above, to the respective cause or effect they were addressing.

Table 8. Summary of Pilot #5b parameters

Pilot #5b – Evaluation	
Number of evaluations per response	3
Amount paid per evaluation	\$0.02

The cause and effect pair were presented, and the Turker was asked to determine if the “events are not related”, as shown in Figure 16 below<sup>1</sup>. The assumption was that if the question posed was difficult to comprehend, the resulting response (Event B) would not be related to the given target effect (Event A).

The image shows a survey question interface. At the top, it asks "What is the likelihood that:" followed by "Event A: S(eventA) will cause Event B: S(eventB)". Below this, there are two radio button options: "Likelihood: %" and "Events are not related". A "Submit" button is located at the bottom of the question area.

Figure 16. Turkers asked to evaluate if Event A and Event B were related

## Results

Average ratings for relatedness were obtained for the different types of question, where “1” corresponds to where the cause and effect pair were related, and “0” if they were not. For questions under the ‘simple cause’ style, the average relatedness was 0.85; for ‘outcome’, it was 0.76; for the ‘OPM’ style, it was 0.60. Given the lower ratings for the OPM style, it is possible that this methodology may have confused the user, or perhaps the raters as well. Regardless, given these results, the eventual prototype platform will be contained to asking for either simple causes or outcomes. Since the type of phrasing within these styles did not appear to matter, this will not be a variable of concern in the eventual platform.

### 3.3.2 Anchoring Bias

The next pilot assessed the variance in the quality of the cause suggested, depending on the extent of other contributions which were visible to the user.

Ansburg & Hill (2003) suggested that the visibility of other answers to a given question can anchor an individual’s subsequent response, resulting in a less creative answer than if no other answers were made available to him. The current pilot therefore evaluated whether the creativity or practicality of a contribution might change depending on whether other responses were visible or not.

Two conditions were used; one where 3 other responses were shown, and one where no other response was shown. The 3 responses in Condition 1 were kept standard for all Turkers given this condition.

---

<sup>1</sup> The figure shows that likelihood was also evaluated using the same task. However this is irrelevant here – the result is used for a separate pilot on curation, which will be described below.

**Propose an Idea**

Propose a way of achieving the scenario that the yellow circle is pointing to.

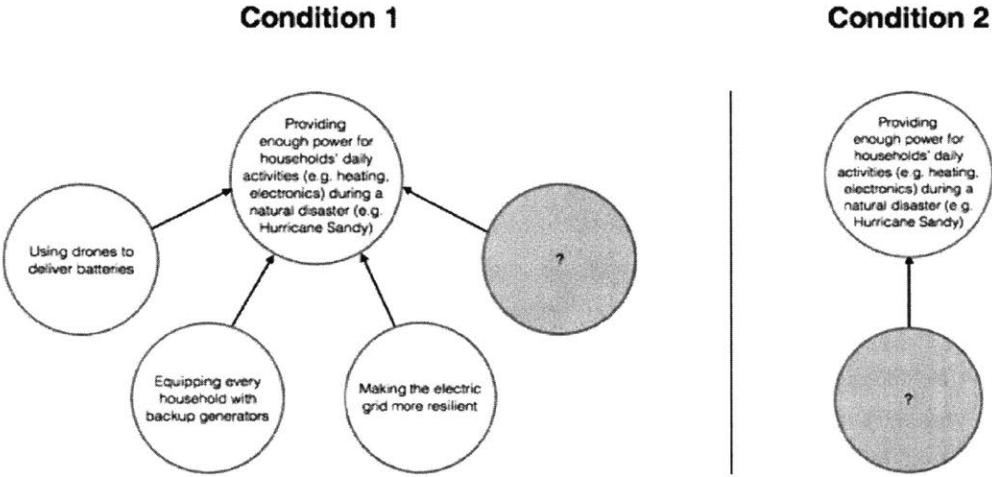


Figure 17. Varying visibility of map to test anchoring bias

Table 9. Summary of Pilot #6a parameters

Pilot #6a – Anchoring	
Number of anchoring types	2
Number of contributors per task	~5
Amount paid per task	\$0.06

**Outcome measure**

A separate Turk task was used to evaluate the creativity and practicality of the causes contributed. Turkers were asked to rate both on a scale of 1-3, ranging from not creative/practical at all, to extremely creative/practical.

Table 10. Summary of Pilot #6b parameters

Pilot #6b – Anchor evaluation	
Number of evaluations per response	3
Amount paid per evaluation	\$0.03



**Survey**

**Problem:** Providing power during a natural disaster

**Solution:** Alternative power sources

1. How **creative** is the solution?

- select one -

2. How **practical** is the solution?

- select one -

**Submit**

Figure 18. Turkers asked to evaluate creativity and practicality of a contribution

## Results

Figure 19 and Figure 20 show the types of responses obtained from each condition, as well as the average creativity and average practicality of all responses in that condition. The number of responses is few (11 in total, since the uptake rate of the task on Mechanical Turk was less than 100% in each), meaning that the ratings of creativity and practicality cannot be meaningfully compared.

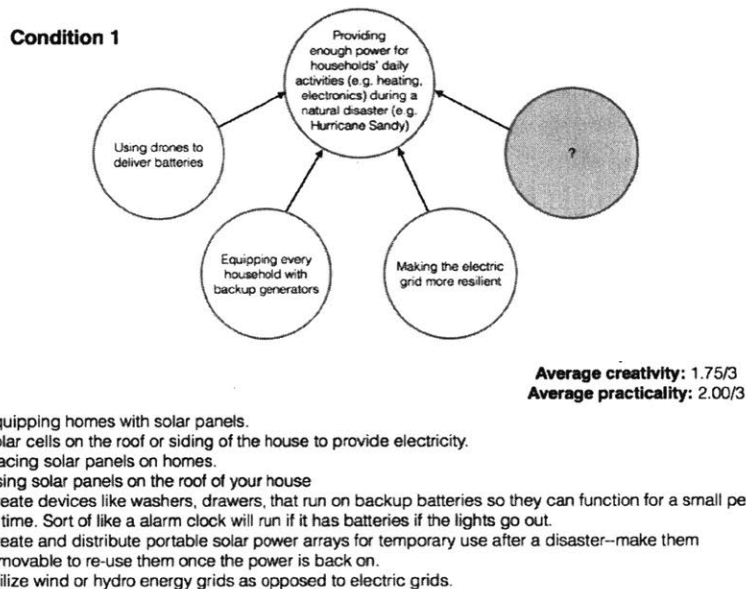



Figure 19. Responses from condition where user is anchored by other contributions



## Condition 2



**Average creativity: 2.00/3**  
**Average practicality: 2.25/3**

- gas powered generators
- Solar Energy
- Portable solar panels
- By using a community wind and solar back up facility each community develops a wind and solar farm spread out in the community using business roof tops.

Figure 20. Responses from condition where no other contributions are shown

The results do indeed correspond with prior research, where the unanchored condition appears to have solutions of higher average creativity. Despite this, one might note that two of these solutions “gas powered generators” and “solar energy” are not particularly creative, and indeed their scores were lower than the third and fourth solutions.

One possible consideration for why noncreative solutions might appear in this non-anchored scenario is the possibility that the anchoring contributions from condition 1 in fact suggested a some of these non-creative solutions, preventing the Turkers from repeating them, and forcing them to think of more unconventional solutions. Perhaps there is a common, uncreative and impractical set of answers to the given question of providing emergency power – or any other question we might decide to yield solutions for – and a strategy to achieve more useful contributions from the crowd might be to in fact display these basic solutions to them.

The takeaway from this pilot is that the effect of anchoring is a nuanced one, and for the sake of the current development of the platform, this variable will not be addressed in detail in this study. It will be taken that the visibility of the map has an indeterminable effect.

### 3.3.3 Availability Bias

An additional pilot was done to determine whether if more of the map was revealed to a user, an availability bias might occur – wherein users would only consider causes close to their current area of focus, versus exhaustively searching through the whole map to find the best place that they could contribute.

Instead of using an actual causal map to determine what areas users might be biased to visually attend to, this question was operationalized in the form of a list instead. Users were asked to provide a cause for a target effect, where they could choose the target effect from a dropdown list of possibilities. In one condition, the dropdown menu was arranged alphabetically; in another, they were arranged randomly.

Table 11. Summary of Pilot #7a parameters

Pilot #7a – Availability	
Types of lists	Alphabetical; random
Number of contributors per list	45 (alphabetical); 86 (random)
Amount paid per task	\$0.02

Choose one of the following...

- select one -

and propose one cause for it:

Submit

Figure 21. Task allowing Turkers to select their own prompt

### Results

In either list case, the prompts at the top of the list appeared much more likely to be responded to than those further down in the list, as shown in Figure 22 and Figure 23. Because this was true in both the alphabetical and random list, it removes the possibility that the best ideas were by chance at the beginning of the alphabet, or at the beginning of the random list.

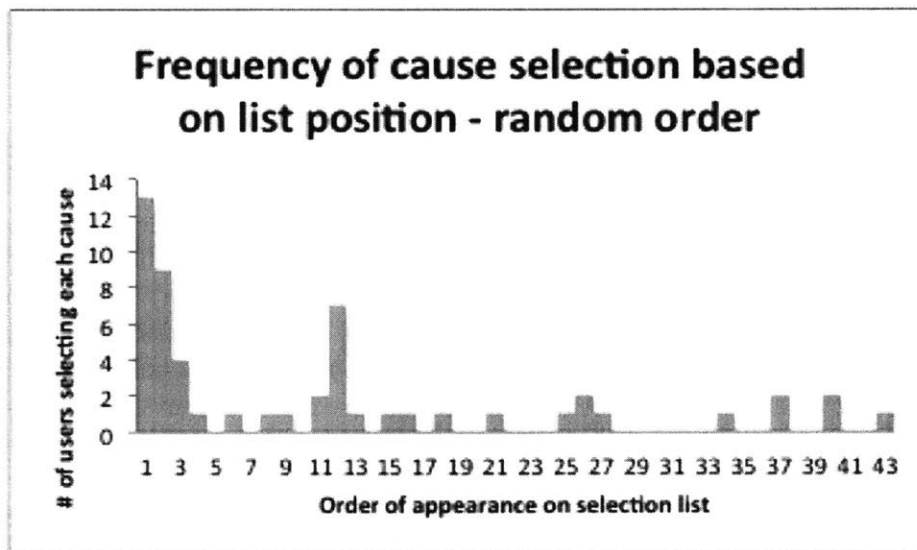


Figure 22. Frequency of cause prompts selected by users, based on order of appearance

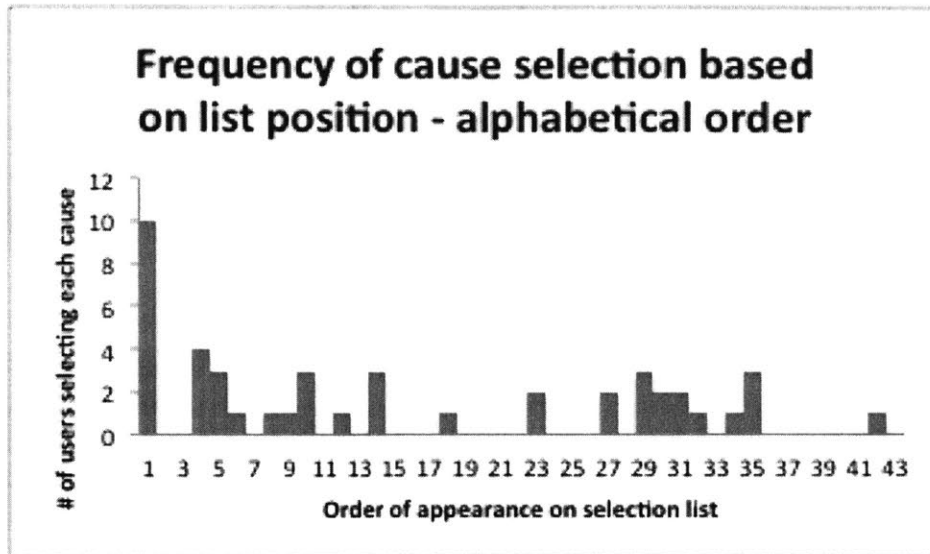


Figure 23. Frequency of cause prompts selected by users, based on order of appearance

What can be concluded from this is that users indeed had an availability bias when looking for causes to attend to, and answered those prompts that were most convenient to access, as opposed to searching through a list to find one they could contribute most usefully to.

This might be taken to mean that given a map, users might be drawn to focus only on certain aspects of it, contributing to whatever might be first presented, as opposed to searching throughout for a useful area to contribute. In order to grow the solution architecture in all directions, as opposed to the ones users may be presented first on the platform, the eventual platform might add features such as randomly presenting parts of the map to the user, as opposed to the whole. As discussed in the prior section, however, the resulting change in anchoring has unclear effects on the quality of contributions.

One important caveat to note here, however, is the amount that Turkers were paid to do this task. One could conceive that being paid very little, Turkers were incentivized to complete the task as quickly as possible, and had no incentive to provide higher quality results. The issue of pay is also a concern in all the other Turk experiments above; however in those experiments, two different cases were tested under the same payment conditions (e.g. testing more or less anchoring, but with the same pay in each scenario) such that the effect of concern might be considered independent of payment. Here, however, the pilot has no null – the question was posed as *whether* they have an availability bias, as opposed to whether they have *more* of an availability bias in one condition versus another.

Designing an exhaustive experiment to test the potential availability bias of users in such a platform is unfortunately beyond the scope of this thesis. However, it is enough to note that the availability bias can be an issue here, and as much as possible, it should be avoided in the design of the platform.

### 3.3.4 Takeaways

This series of pilots have been useful to determine how to pose the question to users so as to solicit the best possible contributions – specifically, what language to use to frame the question,

and what context of contributions to present it in. The takeaways from these pilots are summarized in the table below.

Table 12. Takeaways on how to solicit contributions from users

Proposed Features	Potential Pros	Potential Cons
<b>Simple cause/outcome prompt:</b> Asking users to provide straightforward causes or effects (as opposed to framing it within an specific causality structure)	Able to elicit relevant responses	-
<b>Random presentation:</b> Varying the portions of prior contributions shown to the user	Prevents availability bias Minimize the sensitivity to anchoring	Prevents users from finding the most optimal place to contribute

### 3.4 Curation

The pilots for online and offline collaboration (Pilots #1-4) also suggested the importance of introducing curation, as a means of somehow ensuring that contributions remained simple and easy to understand, and were also in fact causal contributions.

Such curation can incur a large amount of effort, however. It can be as involved as that utilized on OpenIDEO, with official moderators who integrate solutions. A simpler alternative might be a rating system, where individuals rate the comprehensibility or causal accuracy of a given node.

#### 3.4.1 Likelihood as a curator

In the following pilot, likelihood ratings were explored as a means of curating both the causal relationship and the comprehensibility of a given idea at the same time. Users were asked to rate the likelihood of the two connected nodes – that is, the likelihood that one event would cause another. They were also given the option of giving it a likelihood of 0% if the ideas were incomprehensible to them (as was used in the prior section). With such a rating, those ideas which had no relation at all to a connected node (i.e. 0% likelihood) could be easily eliminated.

Using likelihood also allows for a potential value-adding feature in the future, where showing the strength of relationships between nodes in the map might help to illustrate which aspects of the overall solution architecture are better integrated than others. As such, likelihood ratings seemed a useful potential feature to explore here.

#### Building a Causal Map

Given the effectiveness of the simple cause prompts above, the prompt “propose one potential cause for the following:” was used to generate a full causal map from the crowd. Mechanical Turk was used to call for sequential levels of a causal map. Starting with the initial desired outcome of “CO2 levels achieve pre-industrial levels of 280 parts-per-million”, causes were generated by the crowd; then these crowd-generated causes were fed back as the desired outcomes for a subsequent round of tasks (in the “\${event}” field shown in Figure 24. In other words, the children nodes submitted in one task became the parent nodes in a subsequent task.

Propose one potential cause for the following:

S(event)

Submit

Figure 24. Cause generation task given iteratively to Turkers, where each “S(event)” was a contribution from a previous round’s result

This task was conducted 5 times, such that the final causal map was 5 levels deep. At each level, a maximum of 4 children were solicited. If Turkers did not contribute enough responses (i.e. no worker chose to take on the task), then as few as 0 children were obtained.

Table 13. Summary of Pilot #8a parameters

Pilot #8a – Causal map construction	
Number of iterations/map layers	5
Number of contributions solicited per layer	4
Amount paid per contribution	\$0.02

The eventual map included 231 nodes, amounting to 148 full 5-level chains. Samples of the final map are shown in Figure 25 and Figure 26 below.

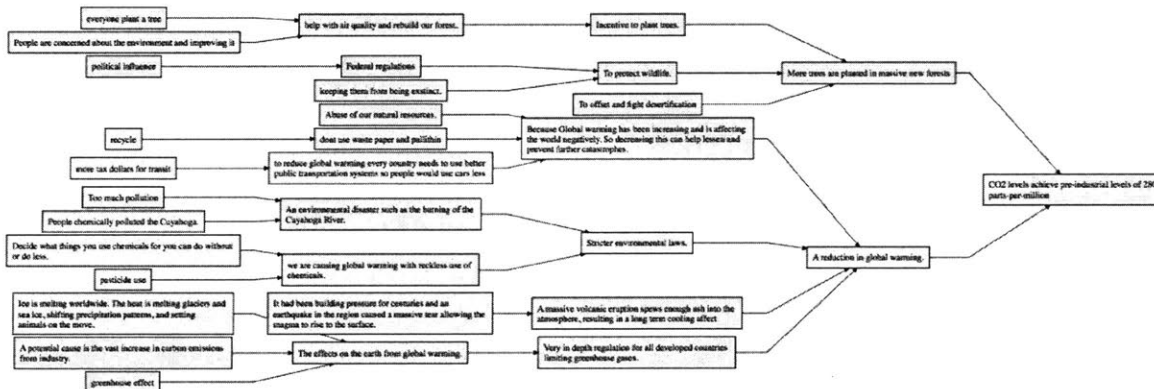


Figure 25. Subset of the map produced by the iterative causal mapping exercise




Figure 26. One of the full 5-level deep chains produced (from political influence to CO2 level reduction)

### Evaluating likelihood

Once this full map was developed, all of the parent-child node pairs were given to Mechanical Turkers, in a separate task, which asked them to assess the likelihood of each.

Table 14. Summary of Pilot #8b parameters

Pilot #8b – Likelihood evaluation	
Number of evaluations per pair	3
Amount paid per evaluation	\$0.02

What is the likelihood that:

Event A: \${eventA}

will cause

Event B: \${eventB}

Likelihood:  %

Figure 27. Likelihood evaluation of cause effect pairs given to Turkers

### Outcome Measure

The outcome measure of interest was the overall creativity and practicality of the whole chain of events. The hope was that if the likelihood ratings could be used to curate out those contributions which had low likelihood, the eventual solution architectures (i.e. the causal chains) were likely to be of greater quality overall.

Table 15. Summary of Pilot #8c parameters

Pilot #8c – Chain evaluation	
Number of evaluations per pair	3
Amount paid per evaluation	\$0.04

The chain of events below is suggested as a way of achieving an outcome.

**CHAIN OF EVENTS:**

\${eventE}  
 which leads to  
 \${eventD}  
 which leads to  
 \${eventC}  
 which leads to  
 \${eventB}

which ultimately leads to the OUTCOME where:  
 USA relying on 100% renewable energy by 2050

Figure 28. Chain of events presented to Turkers

Chains of events were extracted from the causal map, and presented to Turkers as shown in Figure 28 above. Turkers were then asked to rate “on a scale of 1-5, how creative is this chain of

events as a way of achieving the outcome?”, where 5 was “extremely creative”, and 1 was “not creative at all”. The same was done for practicality.

### 3.4.2 Takeaways

From these pilots, it was possible to compile the effect which likelihood ratings had on the overall quality of the causal chain, to see if likelihood could serve as an effective curator for quality. Likelihood could be used as a curator either by having a threshold of average likelihood score (i.e. the average likelihood ratings for each pair across Turkers), or the standard deviation of the score (i.e. the standard deviation across Turker’s likelihood ratings), below which a given node would be excluded.

For example, if the threshold was set at an average likelihood of 30%, then any cause contributed that had an average likelihood score of lower than 30% was excluded from the map (along with all of its children). Or, if the likelihood threshold was set at a standard deviation of 60%, then any cause whose ratings deviated by greater than 60% was excluded from the map (again, with all of its children). The average creativity and practicality of chains in the remaining map was then evaluated, to see if the removal had any effect of the overall quality contributions.

The following graphs show how the overall quality of the map varied as a function of where the threshold was set when using average likelihood as the curator (Figure 29), versus using likelihood standard deviation as the curator (Figure 30).

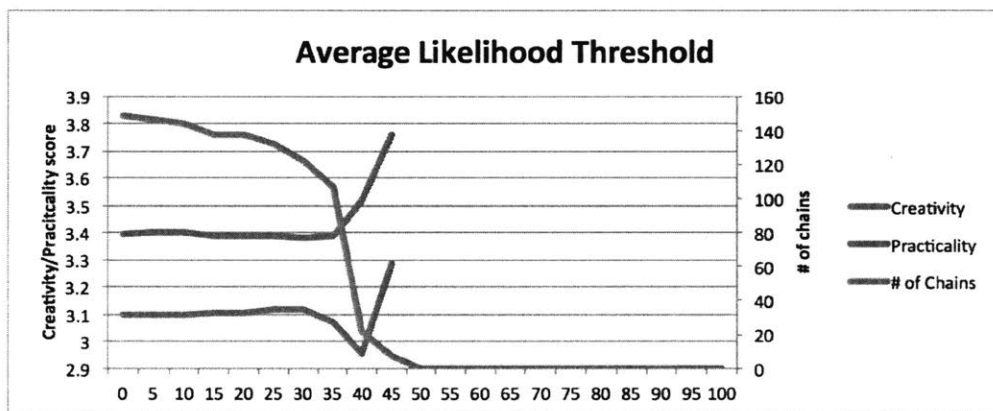


Figure 29. Variation of quality as a function of average likelihood



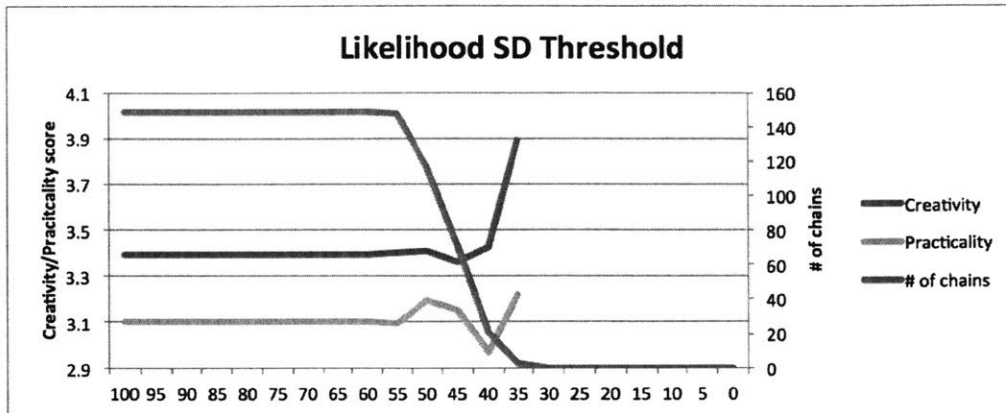


Figure 30. Variation of quality as a function of likelihood standard deviation

The results show that likelihood (both average and standard deviation) can indeed be used as a threshold for curation, but at the significant cost of the number of contributions included. The discontinuity in returns to quality occurs at about the point at which half of the chains have been eliminated. The results also show that there are no significant outliers of low quality contributions that may need to be curated out in the first place.

It is important to note that likelihood can be a counterproductive means of curation. On one hand, it can help to rule out contributions that have no causal relationship whatsoever, which are presumably contributed erroneously or egregiously (of which there appear to be none, or few in this case). However, likelihood is probably correlated with the measures of creativity and practicality, where creative solutions might by definition be less likely, while practical solutions are inherently more likely – such that curation by likelihood may in fact exclude quality solutions. Yet, the results suggest this may not be the case here, given that raising the likelihood thresholds were in fact able to raise both the creativity and practicality outcomes. Thus, at least in the context of developing causal chains for climate change solutions, likelihood is taken as a useful curator.

Table 16. Takeaways on how to implement curation

Proposed Features	Potential Pros	Potential Cons
<b>Likelihood curation:</b> eliminating low quality ideas using likelihood ratings	Able to isolate more creative and practical solutions in this case	May not always correlate with quality



## 4 Collective Causality Platform Design

The pilots in the previous chapter highlighted important features which can be included in the eventual platform design, as well as how they might impact the usefulness and usability of such a platform. In this chapter, these findings on possible features will be extrapolated into a set of design principles, culminating in the selection of a specific design which was developed and used by members of the public.

### 4.1 Summarized Takeaways

The takeaways from the pilots conducted are summarized below, and rearranged to fall into three distinct categories: **attention**, referring to how users choose where to contribute; **contribution**, dealing with how they are guided to contribute when they do; and **moderation**, dealing with how low quality contributions are eliminated.

Table 17. Summary of takeaways from pilots

Proposed Features	Potential Pros	Potential Cons
<b>Attention</b>		
<b>Total map access:</b> Ability to add to any part of the causal map	Allows any part of the solution to be edited	Constrains contributions by anchoring ideas against a very rich solution
<b>Random presentation:</b> Varying the portions of prior contributions shown to the user	Prevents availability bias Minimize the sensitivity to anchoring	Prevents users from finding the most optimal place to contribute
<b>Seed nodes:</b> Pre-populate the map with sample ideas	Encourages simplicity of nodes and commonality of language	Can anchor thinking and reduce diversity of contributions
<b>Contribution</b>		
<b>Leveraging existing methodology vs simple cause/outcome:</b> Educating users in a common modeling methodology (e.g. OPM), vs generic cause/outcome	Reduce confusion, improve creativity	Higher barrier to entry to contributing
<b>Fill in the blanks:</b> Encourage linguistic similarity of nodes	Improves comprehensibility of each contribution	-
<b>Control character limit:</b> Adjust the lower and upper character limits	Promotes simplicity, if bracketed correctly	Potentially excludes valuable ideas with too low of an upper limit, or makes it easier to contribute nonsensical ideas with too low of a lower limit.
<b>Personal refinement:</b> Ability to refine idea separately	Higher quality ideas	Lower volume of ideas
<b>Moderation</b>		

<b>Discussion platform:</b> Some means by which users can clarify or validate others' contributions	Increases quality of ideas	Increases complexity of platform
<b>Feedback mechanism:</b> A way that users can enable prior users to improve their contributions	Improves the quality of each contribution	Can complicate the user's interaction with the platform
<b>Likelihood curation:</b> eliminating low quality ideas using likelihood ratings	Able to isolate more creative and practical solutions in this case	May not always correlate with quality

Some of these features overlap, and are therefore rearranged in the following table to establish clear choices involve in a potential platform design. The check boxes (squares) indicate places where a feature can be considered for inclusion or not, while the radio buttons (circles) indicate where one of the two options must be chosen.

The proposed features were compared based on the likelihood to either enhance the usability of the platform, or the usefulness of its output (i.e. the creativity and practicality of ultimate solutions). These evaluations of impact are based on the findings from the pilots above. It is acknowledged that the results of the pilots can be considered as anecdotal, and that much more research is required to fully understand the extent to which these features might be useful or not, whether for this specific purpose of causal chaining for climate change, or for collective intelligence in general. In some cases, possible features were not tested in pilots (e.g. a possible discussion platform, or feedback mechanism), as such no evaluation can be given as to their potential impact. These fallbacks of the approach of this thesis will be discussed in greater detail in the concluding chapter.

Ratings of either -1, 0 or +1 were given for each feature – where “-1” represents an expectation that it would not be impactful, “0” represents uncertain impact, and “+1” represents the expectation of a positive impact. The filled circles or squares indicate where the features were chosen for eventual inclusion in the prototype platform.

Table 18. Summary of features considered for prototype

Feature		Impactful for usability	Impactful for usefulness
<b>Attention</b>			
<input type="radio"/>	<b>Total map access</b>	0	0
<input checked="" type="radio"/>	<b>Random sample of map</b>	0	+1
<input checked="" type="checkbox"/>	<b>Seed nodes</b>	+1	+1
<b>Contribution</b>			
<input type="radio"/>	<b>Use existing methodology</b>	-1	+1
<input checked="" type="radio"/>	<b>Simple cause/outcome</b>	+1	0
<input checked="" type="checkbox"/>	<b>Fill in the blanks</b>	+1	0
<input checked="" type="checkbox"/>	<b>Control character limit</b>	+1	+1
<input checked="" type="checkbox"/>	<b>Personal refinement</b>	0	+1
<b>Moderation</b>			
<input type="checkbox"/>	<b>Discussion platform</b>	0	+1
<input type="checkbox"/>	<b>Feedback</b>	-1	+1
<input checked="" type="checkbox"/>	<b>Likelihood curation</b>	0	+1

## 4.2 Platform Development

The prototype platform developed was a scenario planning website aimed at forecasting the impact of the US November 2016 elections. The elections were specifically chosen as the context for the website, so as to promote engagement with the website. Following the recommendations of the Collective Intelligence Genome (Malone et al., 2010), one class of features which had not been addressed in detail in the pilots thus far was the “why gene” – the reasons for which individuals would be motivated to participate in the website. The pilots had been conducted on Mechanical Turk, thus “Money” was the motivation to participate in the website. The eventual platform is planned to be a public website however, such that other motivations must be called upon. Thus the election context was used as a means of hopefully activating the “love” gene, where members of the public might be motivated to interact with it for reasons of political interest.

This is quite unlike the original climate change goals of this thesis. However, the timing was opportune to leverage current events to drive users to the platform. It should be noted, however, that all the pilots were conducted in the context of climate change, implying that some of the features proposed there may not be useful for this context. However such differences, if any, were not apparent.

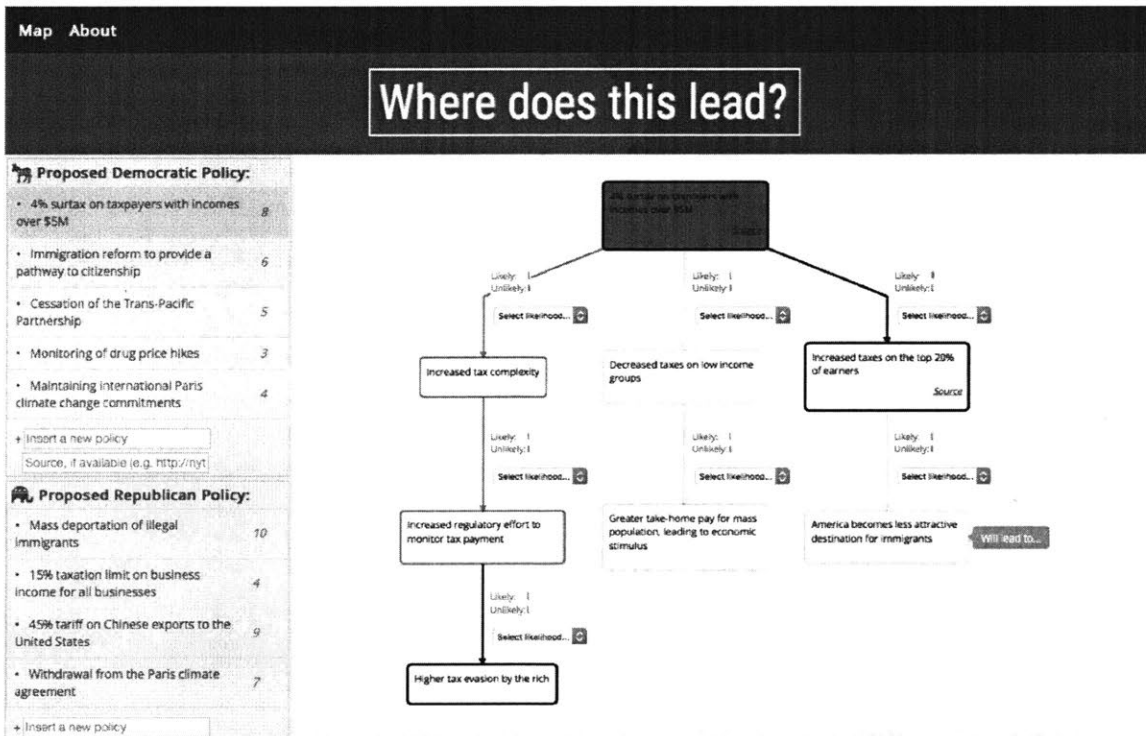


Figure 31. Prototype platform created to test causal mapping in the context of the November 2016 elections

The platform enabled users to consider the downstream impact of the policies proposed by either party leading up to the election. As such, the direction of causality was different from most of that tested in the pilots, where individuals were asked to consider outcomes, as opposed to causes. Through the results of Pilot #5 on question framing, it was expected that users would be able to propose feasible reasonable outcomes.

Users had the ability to add either democratic or republican policies. Each new policy initiated a new map, where the proposed policy formed the root node of the map. Users could then add to the root node, to suggest possible implications of the proposed policy, as well as follow-on implications of those events, and so on and so forth.

### 4.3 Implementation of Design Features

This section describes the ways in which the features proposed from Chapter 3, and selected in section 4.2, were implemented in the prototype platform.

#### 4.3.1 Map Visibility: Random Versus Whole

In the eventual platform, the visibility of the map was not ultimately a choice between a random segment versus the whole – rather, it was a combination of the two. As evinced in Figure 32, some parts of the map were made more opaque than others. The opacity of nodes was linked to the evaluation of their likelihood (as will be detailed below), such that while the whole map was shown, some aspects of it were more visible than others.

An advantage this brings is the ability to allow users to see all aspects of the map, while at the same time keeping them aware of the most ‘obvious’ contributions, which might perhaps direct them to contribute more creative solutions, as was suggested in Pilot #6 (where the user’s awareness of others’ contributions might have caused them to think more widely for alternative solutions).

Map visibility was used as a means of curation as well, where those darker (i.e. more likely) outcomes would draw more attention, and presumably be added to at a higher likelihood than fainter outcomes.

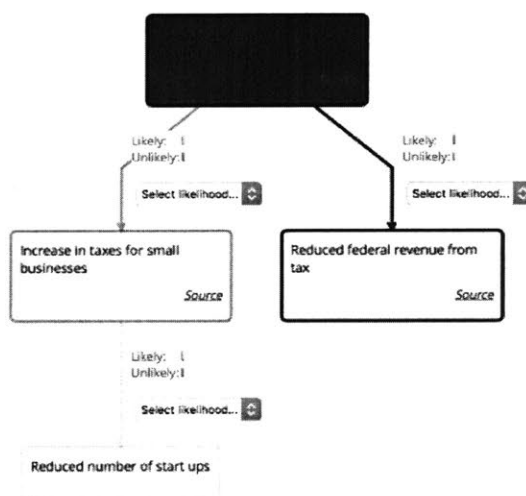


Figure 32. Different opacity of elements of different likelihood in the map

#### 4.3.2 Seed Nodes

The map was seeded by an initial set of policies, as well as outcomes within each set of policies. This was done so as to give users a sense of the type of contributions to be made, and the type of simple language they were to be written in. Indeed, subsequent additions of policy and outcomes

mimicked the style in which these seed nodes were written. Figure 33 outlines in red those aspects of the map which were added before the site was publicized.

One thing of note is that while the seeded contributions included sources (where users could cite a source for where the proposed policy had been stated, or substantiation for why an outcome might be expected), no subsequent contributions included these. This is perhaps because adding sources required a considerable amount of additional effort, and users were perhaps only interested in interacting briefly with the site. Thus, seeds were only able to establish a precedent to an extent.

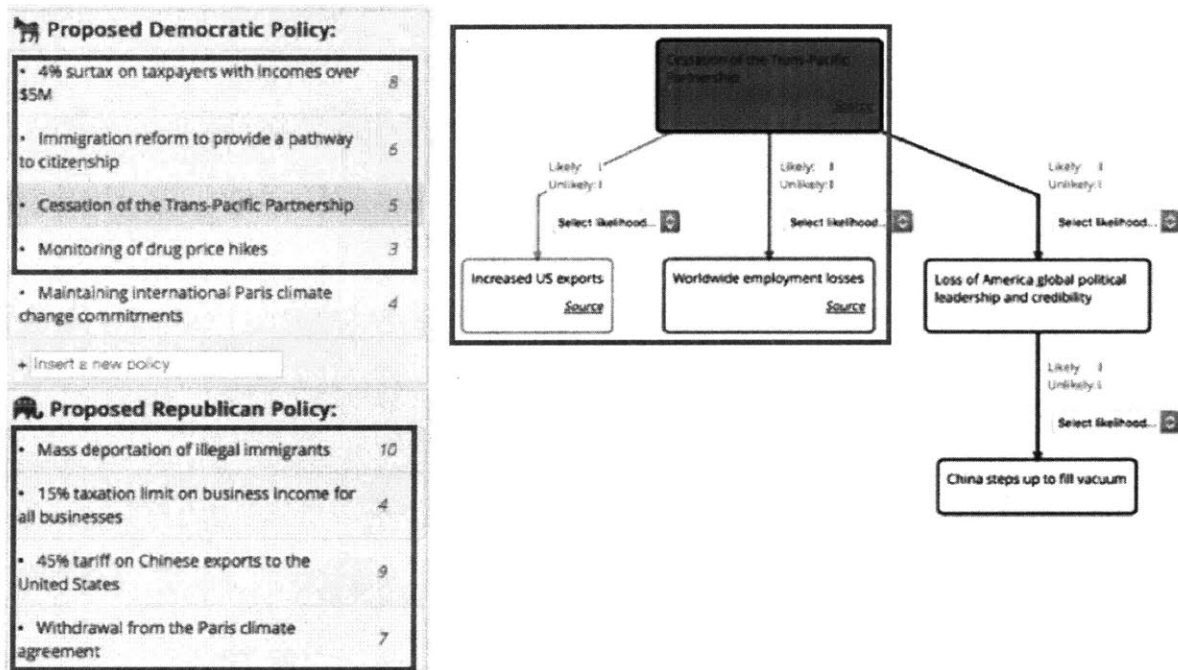


Figure 33. Seeded policies and outcomes created before publication of the website

#### 4.3.3 Simple Outcome and Fill-In-The-Blanks

The prompt to contribute an outcome was phrased very simply. The user was prompted to contribute whenever he moused over an existing node. A “+” mouse would appear, as well as a tag on the node stating “will lead to...”. This prompt suggested to the user that an outcome could be proposed for the given node.

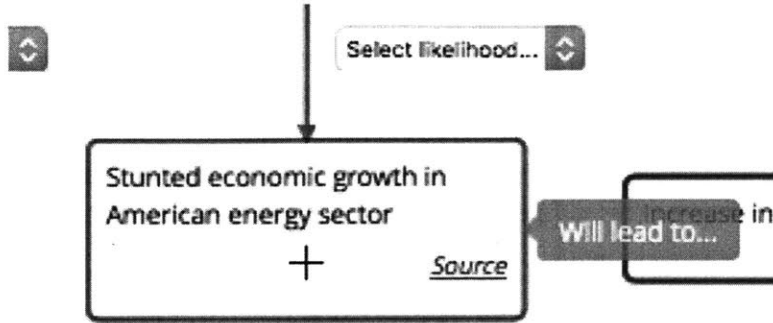


Figure 34. Prompt to contribute an outcome on mouseover of an existing event

Once the user clicked on the node, a simple fill-in-the-blanks prompt, followed by the invitation to “propose a new scenario”, was used to ask the user to propose an outcome, as shown in Figure 35 below.

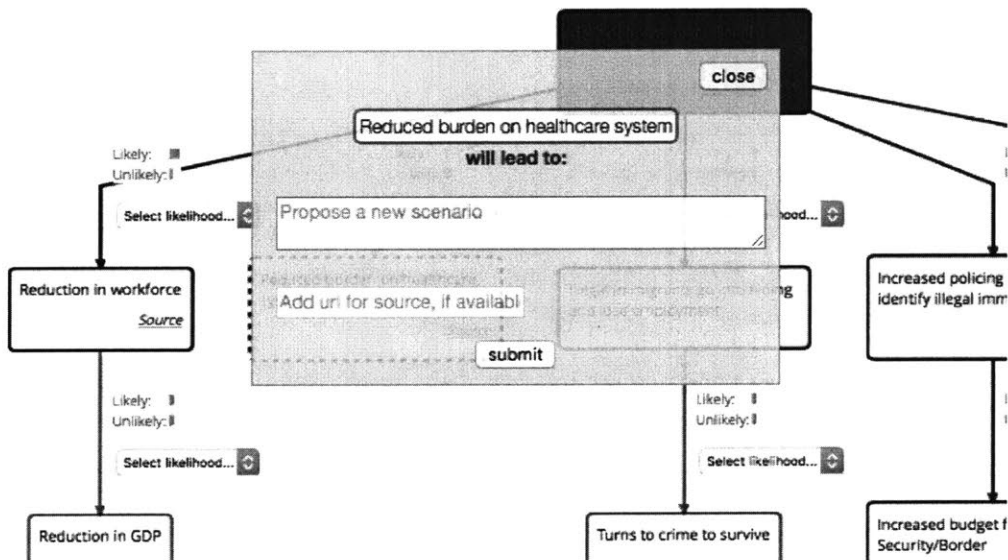


Figure 35. Prompt to fill in an outcome

#### 4.3.4 Control Character Limit

A character limit was set between 10 to 200 characters. As discussed in Section 3.2.2, the ideal character limit is unclear, where too low a lower character limit would possibly increase the likelihood of nonsensical answers being contributed; too low an upper limit might limit quality answers; and too high an upper limit might permit overly complex contributions. This issue has not been resolved here – however, the limits were set by observing a pilot user interact with the website. Her dissatisfaction that many reasonable ideas were too short to be accepted caused the lower limit to reduce from 20 characters to 10.



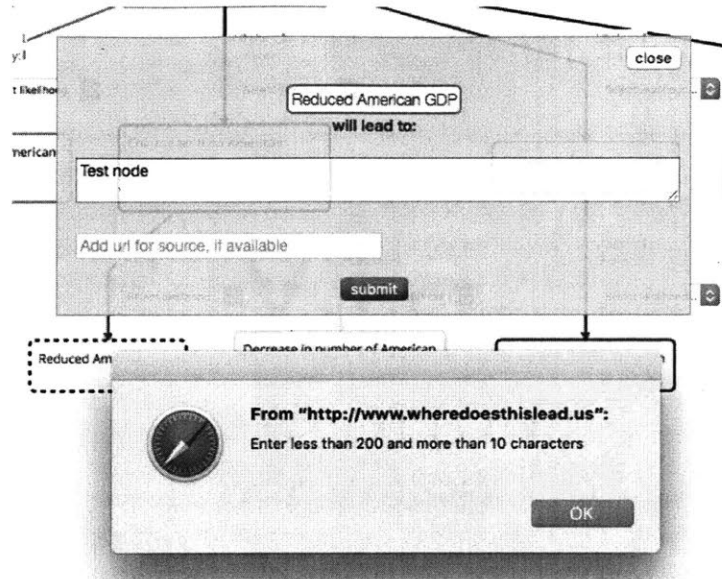


Figure 36. Error alert when submission is not between 10 and 200 character limit

#### 4.3.5 Personal Refinement

The ability to personally refine one's contribution was provided in the form of the possibility of editing or deleting a contribution the user had previously made. No user accounts were required for this website – instead, unique session IDs were assigned every time the user visited the site. This allowed him to have the capacity to edit anything he had added during his visit; but not once he had left the website and come back again. The user was unable to edit other peoples' contributions. The edit and delete interfaces are shown in Figure 37 and Figure 38.

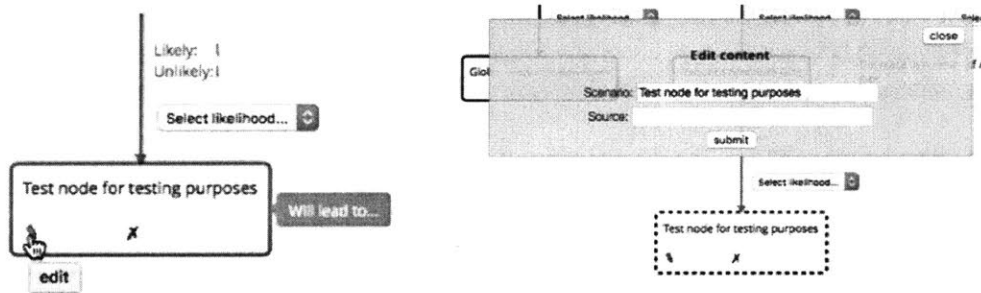


Figure 37. Editing interface

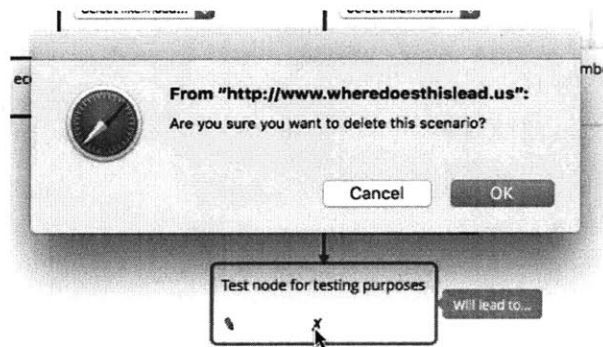


Figure 38. Deleting interface

### 4.3.6 Likelihood Curation

Lastly, as alluded to in the prior section of map visibility, likelihood was included as a form of curation. Along each edge connecting different nodes, users could vote on the likelihood that one node would indeed result in another. This would cause a shift in the small bar chart associated with the edge, showing the proportion of “likely” versus “unlikely” votes which other users had made for that particular edge. In Figure 39, the votes for likely versus unlikely on that edge are fairly even, which is why the bars are of similar length.




Figure 39. Users can vote on the likelihood of each relationship

The opacity of the node also reflects this even voting of likely versus unlikely, where the node is at about 50% opacity. Nodes appear darker then they have higher likelihood ratings for the nodes they point to. Curation is therefore done using visibility, where the (presumably) lower quality aspects of the map are not deleted altogether as was explored in Pilot #8, but instead become less salient to the user.

### 4.4 Platform Usage

All in all, 18 different IP addresses interacted with the site. The site was not published extensively, however even so, the users were fewer than expected. This may be because of interactional challenges with the site itself, where some users reported they were unaware that they were able to change the likelihood ratings on the site. Other explanations are that the level of interest in such an exercise was simply not high.

The following table summarizes the types of interactions done by different users on the site. It is unclear why so few contributed causes as compared to likelihood evaluations; certainly one likely reason is that much more thought and effort was required to contribute a cause. Since the causal chains were only contributed by 3 individuals, it is difficult to assess how creative or practical they ultimately were. Anecdotally, feedback was provided that all of the contributions on the website made sense.

Table 19. Summary of user interactions with prototype platform

User ID	Likelihood ratings	Outcomes proposed	Maps proposed
1	41		
2	29		
3	41		



<b>4</b>	5		
<b>5</b>	16	3	
<b>6</b>	3		
<b>7</b>	8		
<b>8</b>	1		
<b>9</b>	8		
<b>10</b>	8		
<b>11</b>	8		
<b>12</b>	3		
<b>13</b>	7		
<b>14</b>	3		
<b>15</b>	3	5	1
<b>16</b>	6		
<b>17</b>	2		
<b>18</b>	0	14	

## 5 Conclusion

The goal of this thesis was to develop a set of design parameters for developing a collaborative causal mapping platform, which could be used to develop solutions architectures in a manner that online platforms have not done before.

The approach to doing this was largely exploratory, where a series of pilots were conducted to explore the ways in which users interacted when developing causal maps, either online and offline, and the ways in which specific design features might enhance or compromise the quality of users' contributions.

Ultimately, a selection of these features was combined into an actual prototype, which unfortunately had low traction. Therefore results could not be obtained as to how successful such a platform might be at integrating solution architecture.

Regardless, this thesis has helped to highlight features that can affect the success of a collaborative causal mapping platform, and has laid the groundwork for future exploration in this area.

## References

- Amason, A. C. "Distinguishing The Effects Of Functional And Dysfunctional Conflict On Strategic Decision Making: Resolving A Paradox For Top Management Teams." *Academy of Management Journal* 39.1 (1996): 123-48
- Ansburg, Pamela I., and Katherine Hill. "Creative and analytic thinkers differ in their use of attentional resources." *Personality and Individual Differences* 34.7 (2003): 1141-152.
- Cohen, Wesley M., and Daniel A. Levinthal. "Absorptive Capacity: A New Perspective on Learning and Innovation." *Administrative Science Quarterly* 35.1 (1990): 128.
- Crawley, Edward, Bruce Cameron, and Daniel Selva. *System architecture: strategy and product development for complex systems*. Upper Saddle River: Pearson, 2016. Print.
- Diehl, Michael, and Wolfgang Stroebe. "Productivity loss in brainstorming groups: Toward the solution of a riddle." *Journal of Personality and Social Psychology* 53.3 (1987): 497-509.
- Drazin, Robert, Mary Ann Glynn, and Robert K. Kazanjian. "Multilevel Theorizing about Creativity in Organizations: A Sensemaking Perspective." *The Academy of Management Review* 24.2 (1999): 286.
- Dual Use - Off Grid - Illumination Device | InnoCentive Challenge. Jul 17 2007. Web. 09 May 2017. <https://innocentive.com/ar/challenge/5279117>
- Dutton, J. E., and J. M. Dukerich. "Keeping An Eye On The Mirror: Image And Identity In Organizational Adaptation." *Academy of Management Journal* 34.3 (1991): 517-54.
- Henderson, Rebecca M., and Kim B. Clark. "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms." *Administrative Science Quarterly* 35.1 (1990): 9.
- Lakhani, Karim R. "InnoCentive (A)". Case Study. Boston. Harvard Business Publishing, 2008.
- Lakhani, Karim R., Anne-Laure Fayard, Natalia Levina, and Stephanie Healy Pokrywa. "OpenIDEO". Case Study. Boston. Harvard Business Publishing, 2013.
- Lakhani, Karim R., David A. Garvin, and Eric Lonstein. "TopCoder (A): Developing Software through Crowdsourcing". Case Study. Boston. Harvard Business Publishing, 2012.

- Lakhani, Karim R., Hila Lifshitz-Assaf, and Michael Tushman. "Open Innovation and Organizational Boundaries: The Impact of Task Decomposition and Knowledge Distribution on the Locus of Innovation." (May 2, 2012). Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 12-57; Harvard Business School Organizational Behavior Unit Working Paper No. 12-057.
- Levinthal, Daniel A. "Adaptation on Rugged Landscapes." *Management Science* 43.7 (1997): 934-50.
- Malone, T. W., Laubacher, R., & Dellarocas, C. *The Collective Intelligence Genome*, Sloan Management Review, Spring 2010, 51, 3, 21-31.
- Malone, Thomas W., Jeffrey V. Nickerson, Robert J. Laubacher, Laur Hesse Fisher, Patrick De Boer, Yue Han, and W. Ben Towne. "Putting the Pieces Back Together Again." *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW '17* (2017).
- Dahlin, K. B., L. R. Weingart, and P. J. Hinds. "Team Diversity And Information Use." *Academy of Management Journal* 48.6 (2005): 1107-123.
- March, James G. "Exploration and Exploitation in Organizational Learning." *Organization Science* 2.1 (1991): 71-87.
- Narayanan, V.K, and Deborah J. Armstrong. *Causal mapping for research in information technology*. Hershey, Pa: Idea Group, 2005. Print.
- Taylor, C. W. "Some variables functioning in productivity and creativity". In C. W. Taylor (Ed.), *The second (1957) University of Utah Research Conference on the Identification of Creative Scientific Talent*. Salt Lake City: University of Utah Press. (1958).