Distributed Resource Governance Using Asymmetric Anonymity

Research-in-Progress

ABSTRACT

This research proposes a novel method for ensuring fair governance of a common resource using asymmetric anonymity. We propose a system of resource governance that uses anonymous auditors instead of a regulatory agency. Self-enforcement and mutual enforcement are subject to fraud and collusion. In the proposed governance system, anonymity, often associated with negative online interactions, is employed in an asymmetric manner to mitigate these negative effects. This is done by assigning random anonymous auditors to a resource claimant. The claimant's identity is known to the auditor. However, auditors' identities are hidden from both the claimant and other auditors. Cheating, along with improper auditing, will result in penalties for both auditor and claimant. Improper auditing consists not only of allowing unlawful resource use, but also denying lawful use. This paper presents a proof of concept simulation using autonomous software agents and lays out the design for future experiments.

Keywords

Anonymity, distributed system, commons dilemma, governance, autonomous software agents, simulations.

INTRODUCTION

Frequently, by choice or circumstance, the resources of a group are cast into a common pool. An example of a choice would be to pay a periodic fee for insurance to lessen the financial severity of a car accident. A natural resource such a river is an example of a circumstantial shared pool, the use of which is sometimes framed as a commons dilemma (Hardin, 1968), in which the short-term selfish interests of individuals are opposed to long-term group interests. A typical way to ensure proper resource usage is through an enforcing organization. Such organizations commonly incur substantial costs to support staff and facilities to execute rules and laws that govern resource usage.

We propose a type of contract that is a more economical means of achieving proper resource usage by relying on the stakeholders bound to the contract to enforce the rules, rather than relying on a special organization to do so. In other words, it is proposed that a flattened distributed administration might do the enforcement work in a more cost-effective way. It is known that self-enforcement ("the honor system") and mutual enforcement can facilitate cheating and abuse through deal-making, conspiracy and secretive behavior. In the proposed contract, anonymity, often viewed as a bane of online interactions, is employed in an asymmetric manner to reduce side channels used for collaborative cheating and promote honesty by enhancing the suspicion of being observed, thus improving contract enforcement performance. This is done by assigning random anonymous auditors to a resource claimant. The claimant's identity is known to the auditor to promote thorough claim vetting. Cheating and abuse, along with improper auditing, result in penalties. Improper auditing consists not only of allowing unlawful resource use, but also denying lawful use.

Features on the technological landscape that facilitate the proposed contract are the ever-lower latency of online interactions, allowing prompt auditing activity to take place, and the deep reach and rich fidelity of digital information sources, allowing intensive remote observation of resource usage. The proposed contract is developed in two phases. The first phase is a preliminary simulation to better understand the problem space and plausibility of the ideas. The second phase presents a specification of an online game to quantitatively analyze the contract.

A REVIEW OF THE LITERATURE

A Wider Perspective on Distributed Administration

Technology has densely connected people. While we are awash in information, most organizations retain centralized control and authority structures that concentrate power hierarchically. This reflects the divide-and-conquer strategy, with expertise and skilled specialists occupying niches where people collaborate in close contact, featuring high bandwidth, low latency communication.

While an eminently successful scheme, drawbacks of a hierarchical organization include opacity to stakeholders not residing in a particular walled garden, often with concomitant inefficiencies, obsolescences, and malfeasance that take root in less

flexible and more entrenched bureaucracies. People are generally resigned to this eventuality, demonstrated by complacently paying taxes, premiums and dues, dutifully voting and showing up to board, town, and union meetings, all the while ceding authority to duly appointed representatives. That people are often unable to identify their elected representatives, for example, testifies to the enervating and disengaging effect of having information and responsibility without direct authority and involvement.

While not denying the value of full-time experts and specialists, there is also value in allowing diverse organizational stakeholders to exercise some authority in formulating policies, plans and actions from the perspective of those affected by these things (Phillips, et. al., 2009). The customary problem with this approach is that, although stakeholders are motivated for an organization to succeed, they are also outsiders that are neither sufficiently knowledgeable nor available to perform decision-making roles. It is opportune to suggest that technology offers a challenge to this notion. Many complex organizations have already moved in the direction of distributed control with great success -- open source teams are an example of this. The key notion is not just offering openness, but also real authority; without authority interest will wane for many. The advantage is more eyes and hands and less expense. The risk is allowing less knowledgeable people to have a say, and whether they will defer to experts when it is appropriate.

Anonymity

Given the increased use of surveillance technologies in all aspects of human endeavors, it is not surprising that anonymity research has grown in importance and scope. Today, the study of anonymity spans several disparate disciplines. In Computer science research into anonymity has often focused on privacy enhancing technologies like low-latency anonymity networks (e.g., the Tor Project) (Wright & Stepney, 2008). Anonymity systems research typically centers on creating censorship resistant systems, which protect user privacy (Wright & Stepney, 2008).

However, beyond research that explores avenues to improve privacy enhancing technologies, much of the research focuses on the negative aspects of anonymity on human behavior. In economics, anonymity research often centers on the effects of anonymity on decision making, finding that that anonymous contributors are less generous compared to those whose actions and identities are made public (Small & Loewenstein, 2005). In addition, several studies have found that the anonymity of others also affects individual decision-making. Small, Loewenstein, and Slovic (2007) found that participants give less aid to anonymous victims compared to equivalent identifiable victims. Consistent with this "effect of identifiability," Small and Loewenstein, (2005) found that study participants recommended more severe punishments to identifiable wrongdoers than to equivalent, but anonymous, wrongdoers.

Negative Aspects of Anonymity

In the social sciences, much of the anonymity research has centered on the link between anonymity of the Internet and negative behaviors trolling, flaming cyberbullying and cyberstalking. Herring et al. (2002) note that the relative anonymity of online discussion forums provide a new arena for the enactment of power inequities such as those motivated by sexism, racism, and heterosexism. Tokunaga (2010) finds that the anonymity of the Internet allows cyberbullies the opportunity to inflict harm while reducing the threat of being caught. Further, Tokunaga suggests that those who would not otherwise engage in bullying behaviors may do so online due to the anonymity offered by the electronic media.

The link between anonymity and antisocial or unacceptable behavior is not confined to the Internet. Many have suggested that sport fan violence is facilitated by the perceived anonymity of the crowd (Wann et al., 1999). As Rogers and Ketchen (1979) note, anonymity is perhaps the most widely studied disinhibitor of violence.

Positive Aspects of Anonymity

However, other studies suggest that the anonymity of the Internet may produce positive behaviors. In group decision research, anonymous online discussions have been found to reduce the negative effects of power and status differences in decision-making groups (Kiesler and Sproull 1992, Dubrovsky et al. 1991). As a result, Chidambaram (1996) notes that ethnic minorities, women and others who are silenced or sidelined in traditional group discussions often prefer anonymous online discussions over face-to-face deliberations. Wright and Stepney (2008) note that anonymity is widely-used in situations where knowledge of individual users could lead to favoritism, discrimination or collusion (e.g., the marking of exam papers, review of funding applications and the double-blind academic peer-review used for this conference).

RESEARCH METHODS

Phase 1: Preliminary Simulation

A preliminary simulation using autonomous software agents as stakeholders was done to better understand the problem space and plausibility of the concepts. The simulation compared a *regulatory agency* system with a proposed *distributed* one. Stakeholders are enrolled in both systems for comparison. Stakeholders share a pool of commons resources that they can make claims on. In the regulatory agency system, stakeholders pay a fee to administrate claims, ensuring that no cheating occurs. In the distributed governance system, cheating can happen, dampened by the presence of claim auditors.

A run consists of several rounds. In each round, claims are dispersed to a random set of stakeholders. A claim represents an amount that can be lawfully withdrawn from the commons resources. A granted claim increments the resources of the claimant, and fractionally reduces the resources of all the stakeholders. In the regulatory agency system, all claims are lawful and granted, since it is assumed that cheating will be caught by the administration agency. In the distributed governance system, lawful claims are also granted; however, a claim also represents an opportunity to cheat by making a withdrawal for an unlawful amount. Stakeholders have a parameterized tendency to cheat. If unsuccessful, the claimant is penalized by losing his share of commons resources. After each round, stakeholder financial statuses are tabulated.

In the distributed system a parameterized number of auditors are assigned to each claim. Auditors probabilistically decide to allow or deny both lawful and unlawful claims. The resolution of the claim is decided by the consensus of the auditors. If the distributed system works sufficiently well, due to its low overhead, it will tend to pay off better than the regulatory system. However, if cheating is excessively granted, the regulatory system will tend to pay off better for stakeholders other than the successful cheaters. At the end of the run, the performance of each system is measured by the relative number of stakeholders that find it more financially beneficial.

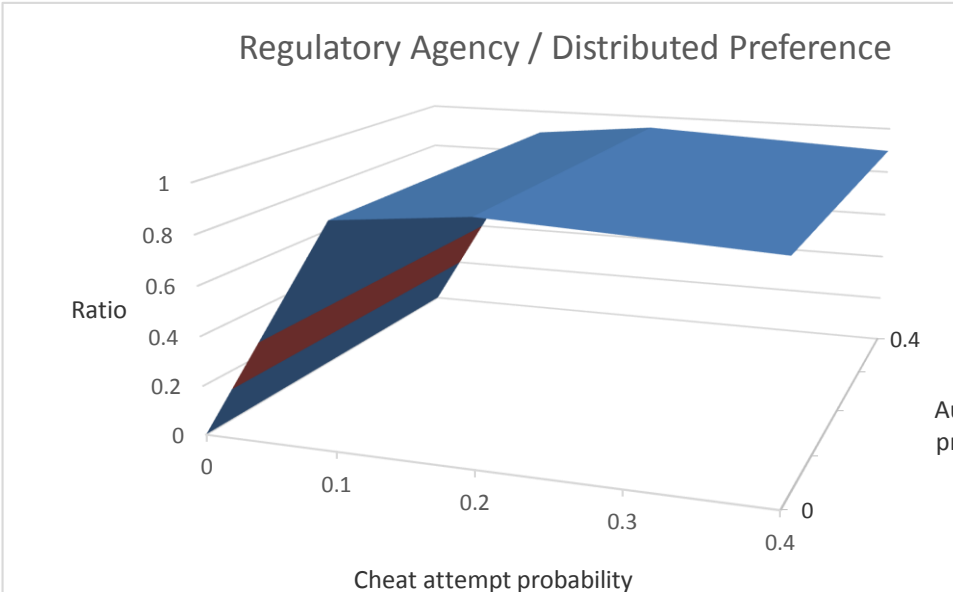

Figure 1 – Regulatory Agency / Distributed system preference ratio with no auditors

Audit error probability	Cheat attempt probability				
	0.0	**0.1**	**0.2**	**0.3**	**0.4**
0.0	0.0	0.10	0.20	0.30	0.40
0.1	0.0	0.89	0.94	0.92	0.89
0.2	0.0	0.89	0.94	0.92	0.89
0.3	0.0	0.89	0.94	0.92	0.89
0.4	0.0	0.89	0.94	0.92	0.89

Table 1 – Regulatory Agency / Distributed system preference with no auditors

Figure 1 and Table 1 are somewhat misleading since there are no claim auditors, but serve as a baseline to show the effect of cheating on stakeholder preference for the regulatory agency system vs. the distributed system. Preference is defined in terms of financial state. With a cheating rate of zero, the distributed system is preferred, since it lacks the overhead administrative fee. However, even a small increase in cheating produces a strong preference for the regulatory agency system, since a majority of stakeholders suffer relatively greater financial loss due to granted cheating claims.

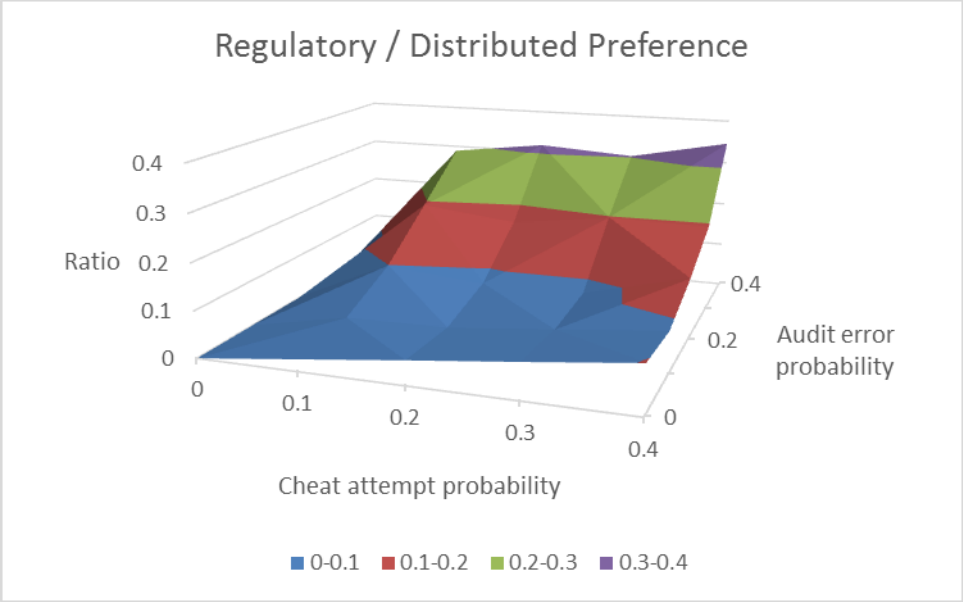

Figure 2 – Regulatory Agency / Distributed system preference ratio with one auditor per claim

Audit error probability	Cheat attempt probability				
	0.0	**0.1**	**0.2**	**0.3**	**0.4**
0.0	0.0	0.02	0.05	0.07	0.10
0.1	0.0	0.04	0.04	0.06	0.09
0.2	0.0	0.09	0.07	0.07	0.13
0.3	0.0	0.19	0.16	0.19	0.19
0.4	0.0	0.28	0.31	0.30	0.34

Table 2 - Regulatory Agency / Distributed system preference with one auditor per claim

Figure 2 and Table 2 show the effect of auditing claims on preference for the regulatory agency vs. the distributed system. When there is no cheating the distributed system is preferred exclusively. As cheating increases, but with errorless auditing to

catch it, the distributed system is strongly preferred. In this situation, it is penalized cheaters that suffer in the Distributed system and thus choose the regulatory agency system. It is only under the conditions of widespread cheating and error-prone auditing, resulting in widespread losses caused by successful cheating, that the regulatory agency system stakes a significant preference, but then only less than 50%.

The main conclusion to be drawn from the simulation is that the presence of claim auditors, even mistake-prone ones, has a powerful dampening effect on the success rate of cheaters, resulting in a significant stakeholder allegiance to the audited distributed system. In Phase 2, more complex conditions involving test subjects will further test the viability of a distributed contract.

Phase 2: Game Experiment

An experiment in the form of an online game will be used to test the performance of a distributed administrative contract using test subjects as stakeholders.

Rules for Experiment

1. Initially, stakeholders each contribute an amount into a resource pool called a commons. This amount is denoted as Rc_i for stakeholder *i*. The commons thus contains $Rc = \sum Rc_i$. Rc_i is always the same value for all stakeholders, signifying an equal ownership of Rc. In addition, each stakeholder has a personal account denoted by Rp_i. This amount can vary in value.
2. A resource entitlement for claimant *i* (Ec_i) is generated from a Gaussian probability distribution P that represents an abstraction of the evidence supporting a resource claim. For example, an entitlement for a pair of shoes would typically run in the tens of dollars, although fringe cases can run down into a few dollars and up into the hundreds or even thousands of dollars. In a real-life scenario, the probability distribution is replaced by substantiations for a resource claim sent to claim auditors. Ec_i is decremented from Rp_i, signifying a loss of resources for the stakeholder.
3. The stakeholder, assuming the role of claimant, uses Ec_i and P to select a claim amount (Cc_i) that will pass auditing and satisfy Ec_i. An amount less than the entitlement might tend to have a better audit outcome, while a greater amount signifies "cheating" that grants resources exceeding the entitlement. Cc_i is selected before any auditors are assigned to the claim and cannot be modified later. Ec_i is not revealed to the auditors as part of the claim.
4. The claim is then assigned a number of random auditors taken from a pool of resource stakeholders. In order to promote thorough claim vetting, the identity of the claimant is known to the auditors, but the auditors remain anonymous to the claimant and each other. The number of auditors is also unknown to all. An auditor is allowed to anonymously communicate with the claimant to provide further claim information.
5. An auditor uses P representing the claim evidence to determine a grant amount (Ga_i). To prevent collaborative cheating, Ga_i cannot be greater than Cc_i. The mean of all the auditor grant amounts ($\overline{Ga_i}$) is the amount of resources granted to the claimant ($Gc_i = \overline{Ga_i}$). Rp_i is then incremented by Gc_i, signifying a resource compensation.
6. To promote claimant honesty, if $Gc_i < Cc_i$, a penalty (Pc_i) is subtracted from the grant that is a function of the difference between the claim and the grant ($Pc_i = f\Delta(Cc_i, Gc_i)$). This will steer claimants away from making excessive claims. Pc_i is then added to the commons.
7. To promote auditor honesty, each auditor is penalized (Pa_i) in proportion to the difference between the auditor's grant and the claim grant ($Pa_i = f\Delta(Ga_i, Gc_i)$). This is meant to discourage both unfair denial and illegitimate generosity to claims. Thus an auditor who colludes with a claimant to grant a large claim runs the risk of penalization by deviating from the grant mean. Pa_i is then added to the commons.
8. There are three ways to score the game:
    a. As in the Phase 1 simulation, a measurement of the preference ratio for a competing regulatory vs. the distributed system can be used as a score. The regulatory agency system exacts a processing fee from all

stakeholders per claim. At the end of a session, each stakeholder will prefer the system that is more financially beneficial.

b. The standard deviation of the stakeholder resources, $R_i = Rp_i + Rc_i$, can be used as a score. A perfect score is zero, indicating that every entitlement/claim/grant triplet was for the same amount and there were thus no penalties. Cheating, denial of resources or penalties will likely skew the resources of the stakeholders; for example, successful cheaters will have relatively more resources than other stakeholders.

c. A measurement of the depletion of the commons resources. This would be computable from the initial and final commons resources and the total entitlements. If the final commons amount is less than initial amount less the total entitlements, then the commons has been excessively depleted.

9. Optionally, since auditor effort has an expense, the resource pool will be reduced for the time spent by auditors to process claims. This will curtail the excessive use of auditors.

Expected Results

It might be expected that the best long-term strategy would be to grant the probability distribution midpoint amount for every claim. However, in the short term, this will result in grants that statistically vary from the entitled amounts. Stakeholders trade transactional privacy in return for the possibility of lower costs. Having stakeholders alternately take on the roles of claimant and auditor encourages cooperative *tit for tat* behavior (Axelrod, 1984).

Independent Variables

As a baseline, the honor system can be implemented and tested by forgoing auditors and sharing this information with claimants. On the other end of the spectrum is an infallible regulatory agency system, which will produce a standard deviation score of zero.

Possible independent variables include:

1. Competing regulatory system claim processing fee.
2. Commons resource amount.
3. Number of stakeholders.
4. Average number of auditors.
5. Penalty functions.

Implementation

The implementation consists of a web game using the Google App Engine (conformativegame.appspot.com) that allows subjects to participate in the experiment. The user interfaces (UIs) for the game have been designed for the three stakeholder roles.

Figure 3 – Home view

Figure 3 shows the home view for a game player. This view allows a player to join (and quit) a game, view his entitled and actual resources, and chat with the other players and the game monitor.

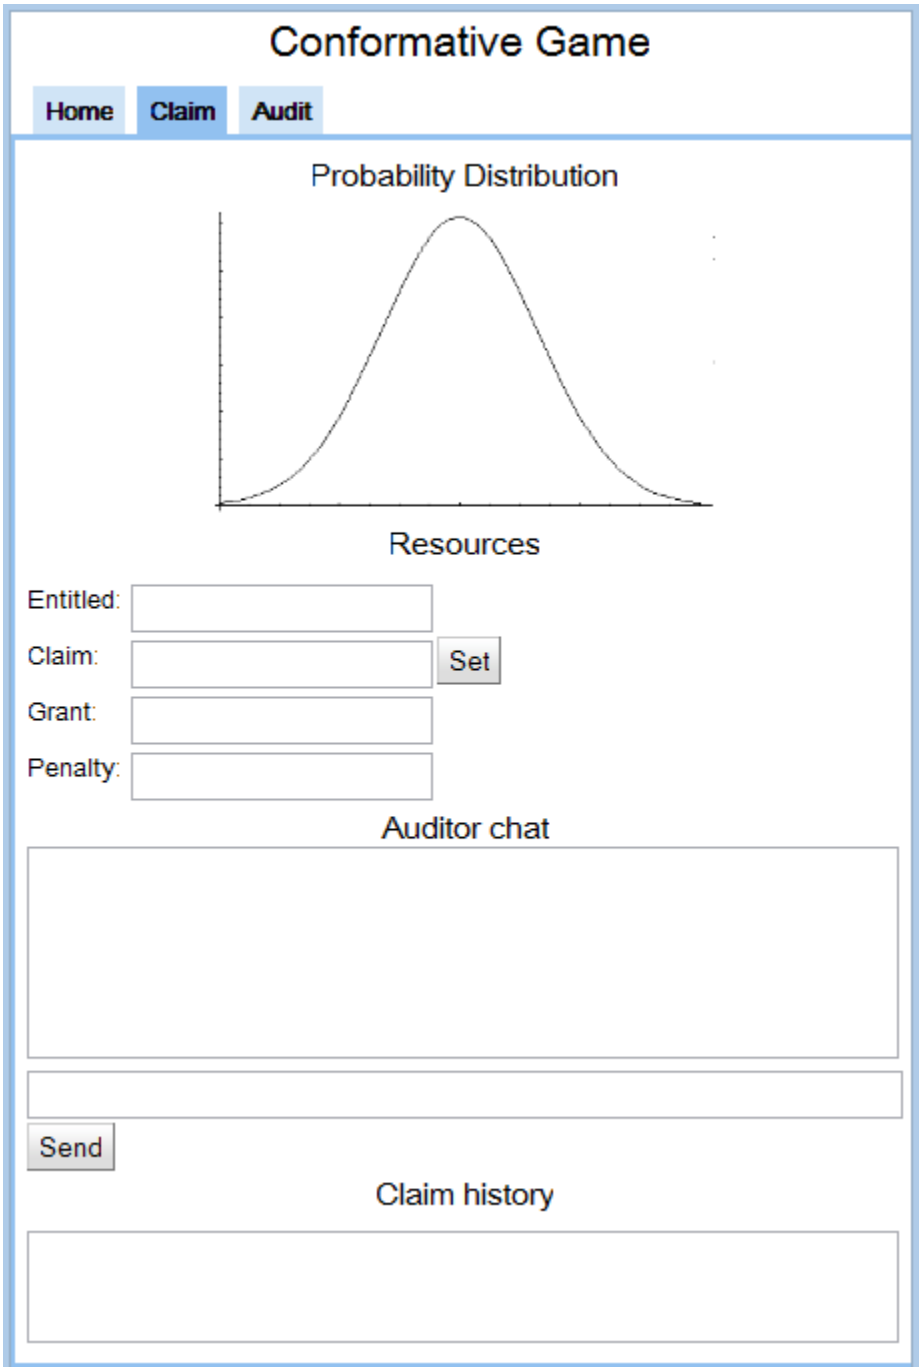

Figure 4 – Claim view

Figure 4 shows the view of a stakeholder as a claimant. The probability distribution from which the entitled resource amount is generated is displayed. The claimant can set a claim amount, then optionally chat with anonymous claim auditors before they determine a grant amount. A penalty is then calculated from the difference of the claim and the grant amount. The claimant can also see a history of his claim resolutions.

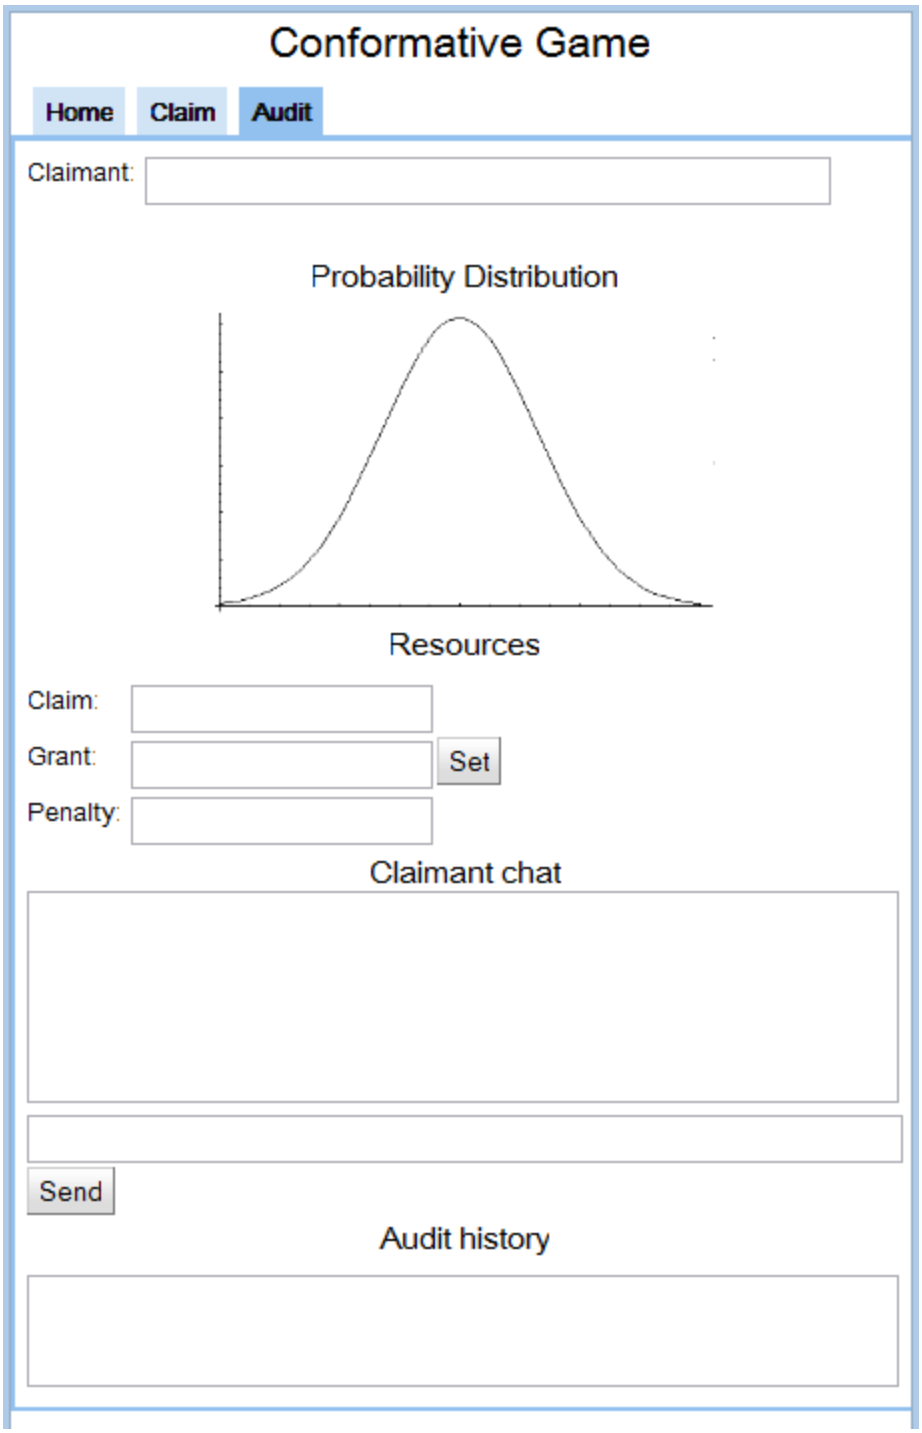

Figure 5 – Audit view

Figure 5 shows the view for the auditor role. The claimant's identity, the probability distribution, and the claimant's claim amount are initially given. The auditor can chat with the claimant to assist in vetting the claim. The auditor then determines and sets a grant amount, from which a penalty is calculated as a function of the difference of the auditor's grant from the mean. The auditor also can view a history of his audit activities.

CONCLUSION

What we hope to accomplish with this research is to demonstrate a workable method for governing a common resource that uses asymmetric anonymity in the role of auditors drawn from a pool of stakeholders. A preliminary simulation points to the plausibility of the method, chiefly by exhibiting a dampening effect that auditing has on cheating. A more thorough test is laid out in the online game experiment with test subjects, the construction of which is underway.

On a larger scale, this study aims to raise awareness of how organizations might leverage a technological landscape that is connecting people into ever denser communication webs. While information sharing has exploded, even to the point of overload, adaptations to improve organizational effectiveness seem to be a patchwork process. Newer teams, producing large-scale open source software for example, have seized on the benefits of high-bandwidth low-latency communications to deploy flatter, more agile, and more inclusive organizational structures. While organizations with differing missions might benefit to a lesser degree, it is an opportunity worthy of consideration.

ACKNOWLEDGMENTS

Provided on acceptance

REFERENCES

1. Axelrod, Robert (1984). The Evolution of Cooperation. Basic Books. ISBN 0-465-02121-2.
2. Chidambaram, L. (1996). Relational development in computer-supported groups. *MIS Quarterly*, 20(2), 143-165.
3. Dubrovsky, V. J, Kiesler S., & Sethna B. N. (1991). The equalization phenomenon: Status effects in computer-mediated and face-to-face decision-making groups. *Human Computer Interaction*, 6(2), 119-46.
4. Hardin, G. (1968). "The Tragedy of the Commons". Science 162 (3859): 1243–1248.
5. Herring, S., Job-Sluder, K., Scheckler, R., & Barab, S. (2002). Searching for Safety Online: Managing "Trolling" in a Feminist Forum. *Information Society*, 18(5), 371-384. doi:10.1080/01972240290108186
6. Kiesler S., & Sproull, L. (1992). Group decision making and communication technology. *Organizational Behavior Human Decision Process*, 52(1), 96–123.
7. Phillips, Katherine W., Katie A. Liljenquist and Margaret A. Neale. 2009. Is the pain worth the gain? The advantages and liabilities of agreeing with socially distinct newcomers. Personality and Social Psychology Bulletin 35: 336-350.
8. Rogers, R. W., & Ketchen, C. M. (1979). Effects of anonymity and arousal on aggression. *The Journal of Psychology*, 102(1), 13-19.
9. Small, D. A., & Loewenstein, G. (2003). Helping a victim or helping the victim: Altruism and identifiability. *Journal of Risk and Uncertainty*, 26(1), 5-16.
10. Small, D. A., & Loewenstein, G. (2005). The devil you know: The effects of identifiability on punishment. *Journal of Behavioral Decision Making*, 18(5), 311-318.
11. Small, D. A., Loewenstein, G., & Slovic, P. (2007). Sympathy and callousness: The impact of deliberative thought on donations to identifiable and statistical victims. *Organizational Behavior and Human Decision Processes*, 102(2), 143-153.
12. Tokunaga, R. S. (2010). Following you home from school: A critical review and synthesis of research on cyberbullying victimization. *Computers in Human Behavior*, 26(3), 277-287.
13. Wann, D. L., Peterson, R. R., Cothran, C., & Dykes, M. (1999). Sport fan aggression and anonymity: The importance of team identification. *Social Behavior and Personality: An International Journal*, 27(6), 597-602.
14. Wright, J., & Stepney, S. (2008, September). Enforcing behaviour with anonymity. In *Proceedings of The Workshop on Applications of Private and Anonymous Communications* (p. 4). ACM.