AVOW

EXPERT-SOURCED FACT CHECKING AND STATEMENT EVALUATION

August 2018

Authors: Nima Bigdely Shamlo and Soheil Alavi
with contributions from many

Abstract

As Thomas Jefferson famously said "An informed citizenry is at the heart of a dynamic democracy". Experts play an important role in a democracy since without being informed by their opinion, citizens cannot make good decisions. But we cannot expect each citizen to carefully study and weight the opinion of different subject matter experts against each other. Hence, we need processes to inform citizens regarding consensus (or lack of consensus) among experts. The problem is that doing so, i.e. establishing an individual's expertise in regards to a particular matter and obtaining opinions from different experts, is a highly non-trivial task: it is often expensive, time consuming and subject to personal bias.

Avow is an expert qualification assignment and opinion aggregation system designed to be scalable, accurate, and trustworthy. Avow enables users to make claims about different statements, then it infers their expertise, using Natural language understanding (NLU) and AIs, on the topics most associated with the statement by considering their education, academic publications, work experience and accuracy of their claim track record. User opinions are then optimally weighted by their expertise and aggregated into a probability value assigned to the statement. These values are then used for fact-checking and to support important decisions.

User will use the platform to establish their expertise in a verifiable manner, similar to how academics use journal publications to do so. Avow utilizes blockchain technology to foster trust and reward users for their contributions with Avow Coins (ERC20 compatible tokens). Block chain also enables users to remain anonymous and data to be provided in a decentralized manner hence fostering freedom of speech and minimizing the likelihood of censorship, data manipulation and censorship. We envision Avow to play an important role in countering disinformation and "Fake News".

Contents

Table of Contents

Acknowledgements

This whitepaper would not have been possible without the contributions of the following individuals: Kamran Bigdeli Shamloo, Tomas Ward and Makoto Miyakoshi.

The Problem

As Thomas Jefferson famously said "An informed citizenry is at the heart of a dynamic democracy". Experts play an important role in a democracy since without being informed by their opinion, citizens cannot make good decisions. But we cannot expect each citizen to carefully study and weight the opinion of different subject matter experts against each other. Hence, we need processes to inform citizens regarding consensus (or lack of consensus) among experts. The problem is that doing so, i.e. establishing an individual's expertise in regards to a particular matter and obtaining opinions from different experts, is a highly non-trivial task: it is often expensive, time consuming and subject to personal bias.

For example, for a political party to advocate for lower or higher taxes to reduce unemployment, there should be a common understanding in the society about the prevalent opinion of economists regarding (a) the current level of unemployment (b) the effect of changing the tax rate on unemployment rate. In the lack of such understanding, there will be no common ground, i.e. a set of agreed-upon facts, to form the basis of a healthy and productive discussion.

For public to gain trust in the opinions of subject matter experts, their qualifications need to be established. Academic and professional certificates are currently the prime indicators of qualification. However, their usefulness is limited when evaluating expertise in regards to specific matters as these matters may lay just outside the person's past experiences or subfield of research. For example, having a doctorate degree in meteorology or physics does not necessarily make an individual an expert in this matter and give her comment a lot of weight on the causes of global warming (but having published a number of highly-cited peer-reviewed articles could). The qualification of each individual for answering a particular question is hence best estimated by taking into account a number of factors, such as their track-record, i.e. how often they have correctly answered to similar questions, their relevant formal education, their employment history, and peer-reviewed publications.

Establishing an individual's qualification in regards to a particular question and obtaining expert opinions are highly non-trivial tasks and currently routinely performed only by a relatively few professionals. For example, journalists often contact academics and other subject-matter experts and ask them to comment on matters raised in the news, e.g. fact-check political or financial statements. This is a highly manual process and has a number of problems:

(a) it is time consuming and expensive to gather relevant information and weight them optimally to establish qualification for a large number of potential experts. Even political fact-checking, a task which may not often require as much research as publishing an article, is currently not scalable to address the demand.

(b) it is time-consuming and expensive to contact a large number of qualified experts and obtain their opinion regarding a particular matter, resulting in a small sample size.

(c) due to the current lack of proven algorithms for optimally weighting different factors, the process could be subjective and hence susceptible to being influenced by ideological biases or

conflicts of interests. This potential for bias can degrade public trust in the outcome and be used by opponents to discredit any conclusions achieved.

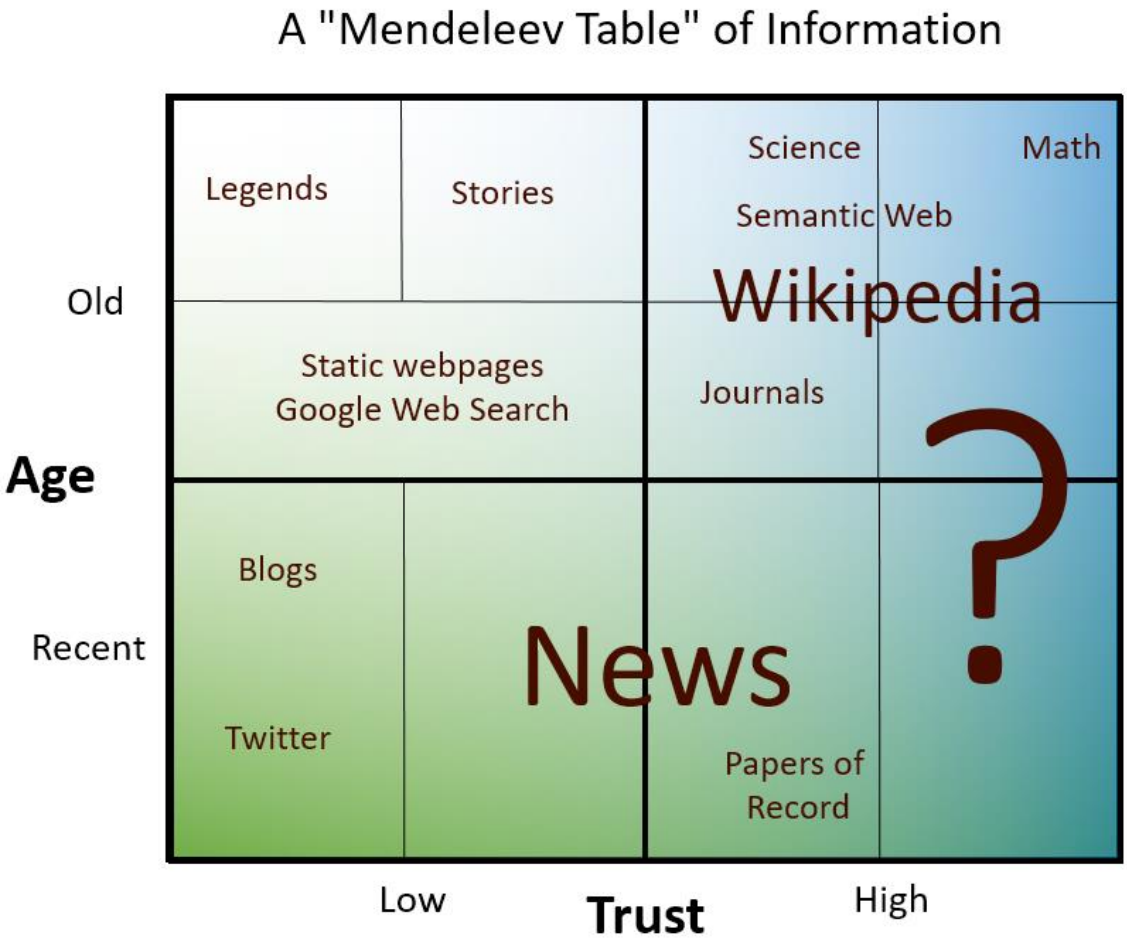

Figure 1. Different sources of information plotted on Age (how old is the information) and Trust (how trustworthy the information is) axes. There is an empty corner for recent, highly trustworthy information.

Figure. 1 shows different sources of information plotted on Age (how old is the information) and Trust (how trustworthy the information is) axes. Many decisions require information that are both recent and trustworthy, however, there is currently a gap in this corner.

The Solution

Here are the properties of an ideal expert qualification assignment and opinion aggregation system:

- **Quantifiable:** it should be possible to calculate the accuracy of system conclusions, e.g. by comparing them to established outcomes (which are obtained in the future), or its own future output when more information has been made available.

- **Accurate**
- **Trustworthy:** by being transparent and verifiable.
- **Fast:** results should be produced in the appropriate time-scale, e.g. before ballots are casted in an election.
- **Scalable:** both in throughput, i.e. number of questions, etc processed and the number of factors considered in establishing relevant qualifications for each individual.
- **Cost-effective**

In addition, users of the system should benefit fairly from their contributions. An ideal system would provide the following benefits to its users:

- They should be able to use the outcomes from the system.
- They should have a degree of "ownership" regarding their contributions, for example, they should be able to download all their contributions in a format that makes it usable in similar systems. They should also financially benefit from the monetary transactions on the system to the degree on which these transactions are based on their contributions. The contributors should also have assurances for safeguarding their privacy, freedom of speech and not tampering with their track record.
- Users should be able to benefit from their reputation and their track-record in other context, such as when seeking employment (the way academics use their publication records to obtain academic positions).

Avow is our proposed system to address these requirements. It is designed to jointly maximize desired properties mentioned above.

The Design

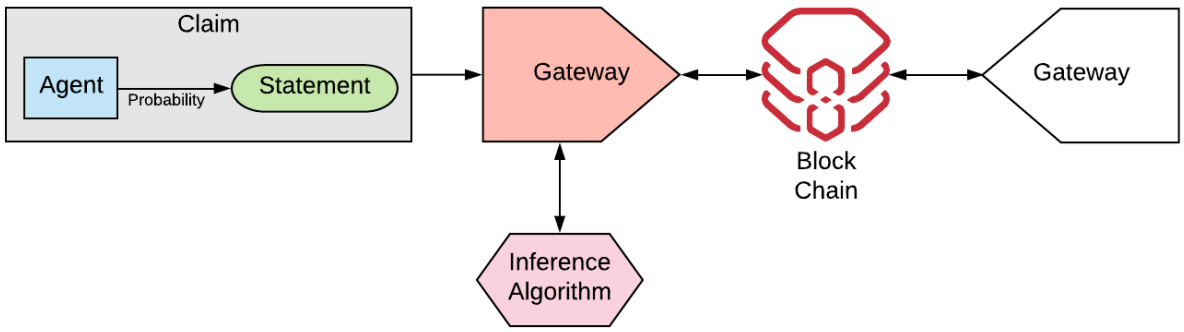

Figure 2. The main components of Avow system.

Avow is composed of five main components:

1. **Agents**: these are users of the system such as individuals, institutions and AI algorithms
2. **Claims**: these are records that hold information regarding "Who Said What, indicated as (Agent, Statement, Probability):
    a. **Statement**: a clear and concise piece of text written in a way that is well understood, e.g. by Natural Language Understanding (NLU) systems.
    b. **Probability** captures the likelihood that the Statement is true.
3. **Distributed database**: containing claims and smart contracts (similar to Ethereum, but with more scalability)
4. **Inference Algorithms**: these infer topic-dependent qualifications and aggregate agent claims to obtain to their best estimates of statement correctness probabilities, e.g. by performing a weighted average. To evaluate agent expertise in relation to each claim, the claim statement will be mapped into a number of topics. For example, a statement about global warming will be associated strongly to "climate science" topic and less strongly into "Physics" and "Science" topics. Each agent qualification will also be scored for different topics. Agents will be deemed more qualified to make a claim about a statement if they score high in topics that are most associated with the statement.
5. **Gateways**: provide user-friendly interfaces to agents and connect them to other components of the system. For example, Avow website will be gateway to the system. Anyone will be able to create a gateway. Gateways obtain a share of the revenue created from the claims managed by them.

The Technology

In the following we explain the design choices and relevant technologies for each component:

Agents

Each individual may be presented by a number of agents in the system. This provides a degree of privacy as different agents may be used to make claim about sensitive issues, hence avoiding punitive actions by oppressive governments, etc. The system will use cryptographic account similar to Bitcoin and Ethereum Accounts. Automated fact-checking algorithms will post their finding as Agent.

Claims

A claim represents the smallest unit of verifiable information: each claim statement is theoretically either correct or incorrect (in the practice we use a probability value since in general we cannot assume certainty). For example, a news article could be represented by a series of claims and the correctness of each to be estimated. The use of concise statements and scalar probability values enables the application of Natural Language Understanding (NLU) algorithms which are needed to make the system scalable. Each claim that enters the system will be cryptographically timestamped (anchored to a public Blockchain such as Bitcoin), providing a verifiable proof that it was made by the agent at a particular time.

Distributed Database

A distributed database implemented using cryptographic principles (similar to Bitcoin or Ethereum) will provide trust and resiliency. Data will be distributed (copied across multiple nodes) and not controlled by any central authority. This makes it extremely hard to censor, or in some other may, manipulate the data. It is also impossible for an agent to deny that they have made a particular claim at a certain time. This is important since the accuracy of agent claims across time will be a strong signal in inferring their qualification. Avow will use blockchain technologies that we have identified to enable a high level of throughput and scalability while maintaining desired decentralized properties similar to Bitcoin and Ethereum.

Inference algorithms

Avow will support a number of inference algorithms (IAs). These algorithms may be developed and maintained by different entities. IAs will compute topic-dependent qualification for agents and aggregate agent claims to obtain estimates of statement correctness probabilities. Here is an example on how an inference algorithm could infer claim probabilities using weighted averaging:

1. C: array of claims $c_i \epsilon C$
2. Each agent A_a to assign the probability $0 \leq p_a^i \leq 1$ to claim c_i
3. For each agent A_a assign the qualification score $0 \leq q_a^i$ depending on the topics associated with claim c_i and the qualifications/expertise of the agent in these topics.
4. Infer the probability of claim c_i: $p_i = \frac{\sum_a q_a^i p_a^i}{\sum_a q_a^i}$

We anticipate IAs to become increasingly sophisticated as more data becomes available in the system. In early stages of the system IAs may heavily rely on public information such as LinkedIn profile data and academic publication records for assigning agent qualifications. In later stages, IAs will be able to utilize the claim history of each agent and relationships across claims containing different statement, e.g. claims made by agents about the qualifications of other agents.

Avow does not assume any particular source of truth. Each IAs may utilize changes in claim probabilities across time as an additional signal in assessing agent qualification since as time passes, often more information about a statement becomes available. IAs may also use external sources of information such as public databases and Linked Data (RDF Triples).

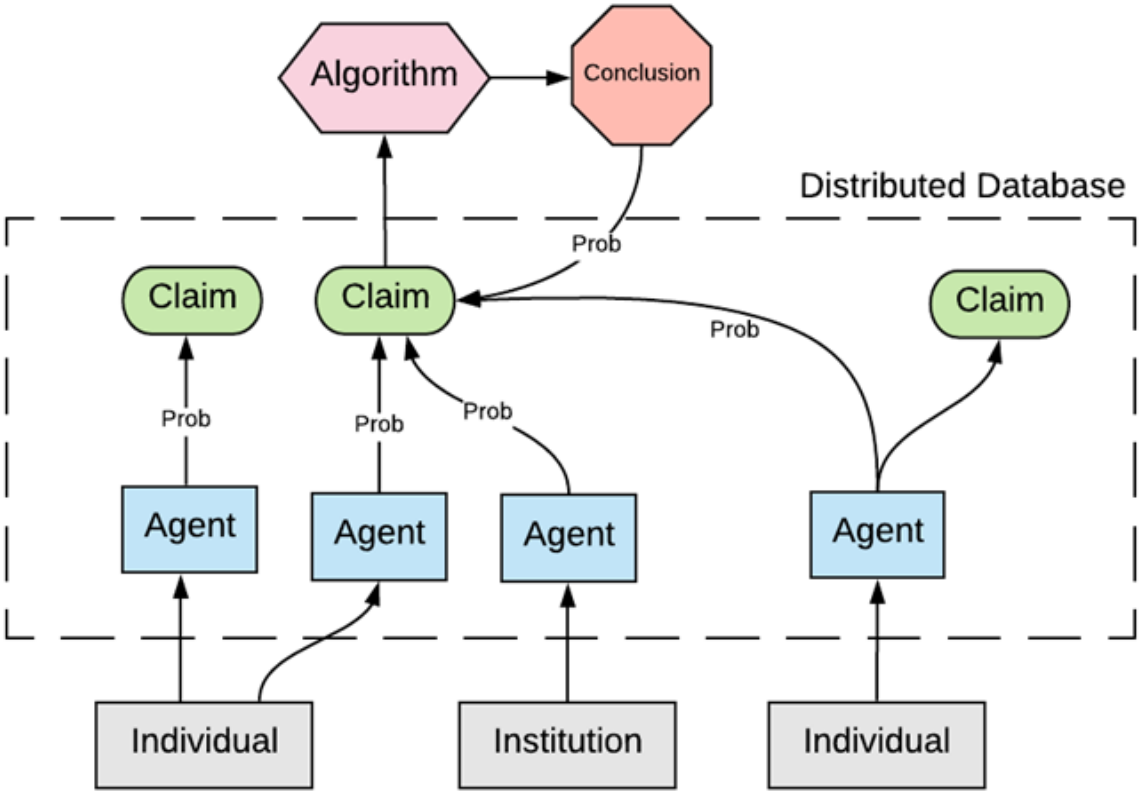

Figure 3. The relationships between Agents, Claims and Algorithms.

To foster trust and enable the evaluation of algorithm quality, IA inferences (aggregated statement probabilities) themselves should be ideally entered into the distributed database as claims made by the IAs. This will also enable the creation of a hierarchy of IA claims based on claims from other IAs (and assigning topic-dependent qualifications to IA by other IAs, similar to what is done for agents), see Fig. 3. Open-source algorithms can be used at the top level to rank closed-source proprietary inference algorithms, see Fig. 4.

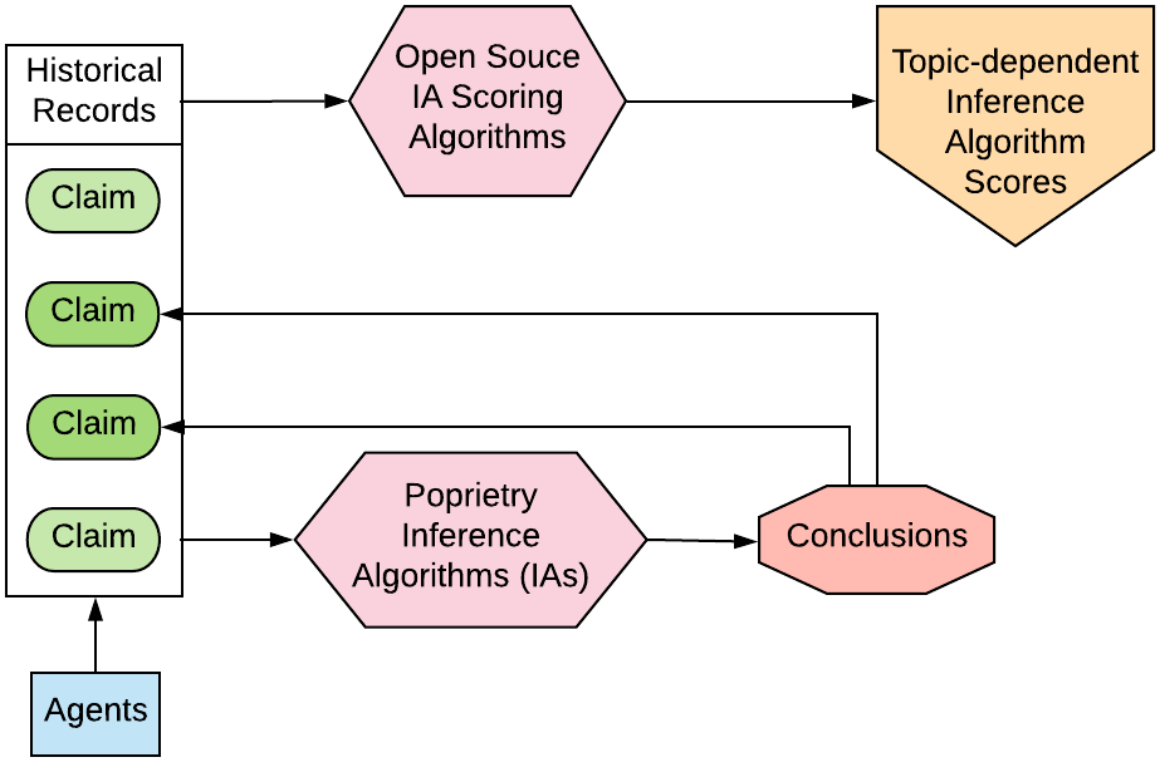

Figure 4. Using open source inference algorithms (IAs) to rank other IAs.

Gateways

A gateway provides a user-friendly interface to agents and connects them to other components of the system. Anyone will be able to create a gateway. Gateways obtain a share of the revenue created from the claims managed by them. Avow website/mobile app will be the first gateway to be created. It will enable users to:

1. Create new claims. They will be told that claim is being placed in the blockchain and can never be deleted. The interface will help to make claim statements short and clear (e.g. using natural language processing)
2. View other users' claims and create claims based on them (assign a probability to the statement made in another person's claim, hence creating a new claim).
3. Search claims by keyword, topic (politics, sports), time (recent), controversiality (many different probabilities assigned to it), etc.
4. Users can ask their friends and social network to make claims, for or against the claims they share with them (by sending a link).
5. See the network of agents (other users, AIs...) that have assigned probabilities to each claim statement. This includes the output of different algorithms (or just one algorithm at first) regarding the aggregated probability that the statement is correct. The algorithms may simply average the probabilities (error prone) or use qualifications (once enough data becomes available) to weigh different claims about the same statement.

6. View qualifications, related to each claim or claim topic, assigned to them and others by inference algorithms based on their qualifications and claim history.
7. Create new "agents". Each user can create multiple agent and make claims through them (to protect their privacy, etc.). The system will prevent multiple agents that belong to the same user to make claims regarding the same statement.

More advanced features:

1. Viewers of claims can encourage agents who provide high quality claims by donating AVW tokens.
2. Provide a list of references, e.g. URLs, when making a claim. References are resources that support their claim.
3. Provide a list of "inspirations". These are the resources that inspired them to make the claim. For example, they point to a news article that says that some politician mentioned that global warming is a hoax, which inspired the user to make a claim in our system and link it to that article as inspiration.
4. An API that provides all claims inspired by a URL. This will allow a newspaper, blog, etc. to display claims from our system related to their article on their page (e.g. along with their inferred probabilities).
5. The system will suggest to the user, based on their qualifications, statements on which they should make claim, e.g. suggesting climate scientists to make claims on the statement "Global warming is a hoax".
6. Allow users, e.g. journalists, to find users that are qualified to make claims for a given statement (may not have been made into claim yet, e.g. what the journalist is investigating). The results could be public claims or private claims (just to the journalist). Private claims may become public automatically after a "sanction" period. Parts of this service would be premium and involve exchange of AVW tokens between the journalist, Avow website and the qualified users.
7. The website will eventually include dedicated interfaces for premium users such as journalists and investors who seek expert opinion (in form of claims, phone conversations, etc.).
8. An agent can "defer" to another agent. This means that they will claim that the other agent is highly qualified to make a claim regarding the topic. If the agent in future turned out to be incorrect about that claim, this will reduce the qualification of the user who has deferred to that agent. A user may defer to a number of agents.

Fig. 5 shows an example of Avow website user interface.

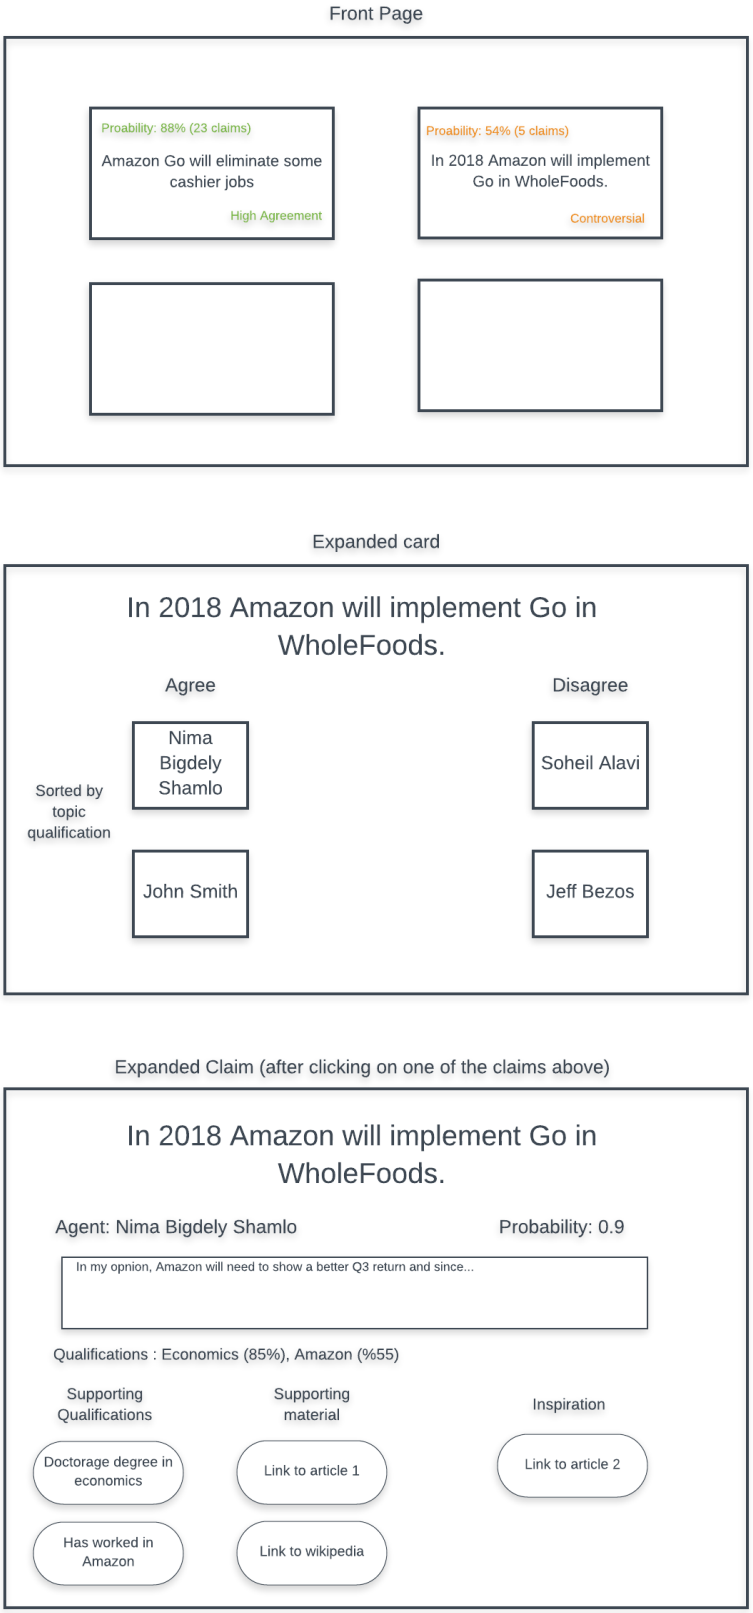

Fig.5 A mockup of Avow website interface.

The front page will contain a series of information cards (Fig.5 top), each associated with a particular statement and displaying its expertise-weighted probability. The presentation of these card will be customized based on user expertise, preferences and topics trending in their geographical region. User can expand each information card to see an overview of claims made for or against each statement (Fig.5 middle). They can add their own claim about the statement or click on any of the presented claims to further investigate each by reading its comments, looking at the qualifications of the agent making the claim and visiting supporting links included in the claim (Fig.5 bottom).

The Benefit to Users

An important question is what service does Avow will provide to users? Why should they return to Avow website? These are the primary benefits that Avow will provide to its users:

- For users who post claims: to establish a track record which shows their expertise in a number of areas. For example, by being among the first to have correctly predicted the result of an election or sports game on a number of occasions indicates that the user has a good grasp of that area of politics or sports. Avow will (a) records these claims in a verifiable way in the blockchain (b) considers users track history *in an automated manner* among many topics. User can then benefit from this establish expertise in a number of ways, e.g. by attracting more followers on social media, by answering questions on platforms that provide monetary reward, or by even in their job applications.
- For visitors to Avow website: to cut through the clutter and clearly see summaries of expert opinions.

The Revenue Model

Direct revenue sources:

- There could be a fee for accessing claim probability values. Claim statement and metadata will be public so the relevant claims can be discovered and paid for. The fee for accessing claim data will be set beforehand.
- There will also be a fee for accessing the result of inference algorithms. If these have been put in the system as claims then the mechanism above will cover such transactions. Social Media and news aggregators, notable Facebook and Google who seek a solution for their "Fake News" problem would be the prime customers for this service. Facebook has already started experimenting with crowd-sourced fact checking (link). Avow offers a more fundamental solution to this problem.
- Agents may also request other agents to provide claims regarding a particular statement in the specified time window and offer a set price for such claims.
- Gateways generate revenue by performing the above transactions of behalf of agents and retaining a portion of the revenue. Even though individual agents may use the system directly

(similar to how one can access Ethereum via local wallets), we anticipate that almost all transactions to be performed via Gateways.

- There may be advertisement placed on a gateway website, but this advertisement can only be paid for using AVW token, e.g. via LinkExchange.io. This incentivizes users to obtain AVW token since it could then be sold to advertisers.

Indirect revenue sources (revenues generated by inference algorithm providers):

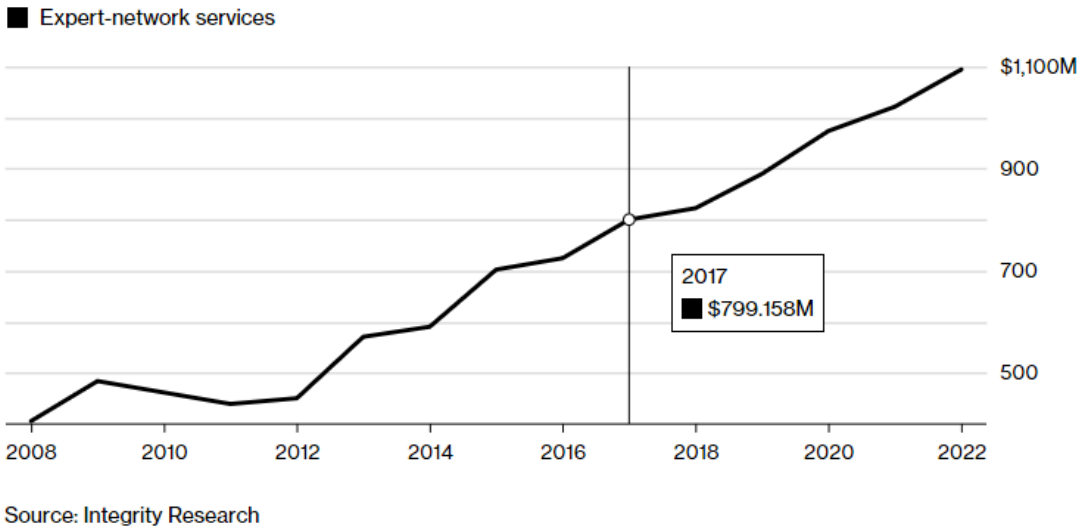

Figure 6. The size of Expert-network services market.

- Expert discovery: there are a number of situations where finding experts who are best qualified to answer a certain question, or comment or a statement is quite valuable. These include "Expert Network Services": Investors seeking answers that will help them to make better investments. This is currently a $800 million dollar market and is expected to grow into a $1.1 B market by 2020 (see Fig 6).
  Topic-dependent qualifications computed by algorithms operating on Avow platform will directly serve this application.
- News: in their book The Elements of Journalism, Bill Kovach and Tom Rosenstiel identify the essence of journalism as "a discipline of verification" and a its first obligation to the truth. Avow is highly inspired by journalism and will strive to serve journalists in their endeavour:
  - Trust in traditional news organizations is at historic low (source 1, source 2). This trust can be partially restored by providing more transparency to the journalistic process, for example the network of sources for an article may be provided alongside each article. Readers will be able to investigate the qualifications of these source. Even anonymous sources can be evaluated looking at their claim track-records (and by claims made by other sources in regards to their reliability).
  - By making the process of finding and obtaining information from qualified experts faster, cheaper and more accurate, Avow will increase the quality of news articles and allow journalist to focus on other aspects of their work such as coming up with important questions to ask.

Even though newspapers have been experiencing a decline in total revenue, the size of this market is still large (was over $14 billion dollars in 2015).

- Recruitment services: journal publications have been the prime method of establishing qualification for obtaining academics positions for decade. Recently, contributions to open-source projects, e.g. GitHub code commit history, are regularly used as important factors in determining qualification in the software development business (link). Avow will enable individuals in other professions to also benefit from their claim track record.

Coin Economy

All payment transactions on Avow platform will use Avow cryptocurrency-based tokens (ticker symbol: AVW), a fungible and tradable asset. Users will be able to purchase and sell AVW tokens (coins) in exchange for fiat currency or cryptocurrencies (Bitcoin, Ethereum).

Initial Coin Offering

All AVW tokens are "utility tokens" and will be pre-mined: a set number of AVW token will be minted using smart contracts at the token creation event and no further tokens will be created afterwards (enforced by smart contracts). As the demand for these token increases, e.g. entities spend these tokens to perform actions in the system, their traded value will increase, providing revenue a mechanism for token holder.

Depending on applicable SEC regulation, part of the funds for the development of the system may be raised via initial coin offering (ICO).

The Control Structure

The early stages of the development Avow will be developed by a commercial entity. Once the system reaches a certain level of maturity and has a diverse community of stakeholders (users, gateways...), the control of the core system elements will be transferred to a foundation. Gateways will continue to be operated independently of foundation. The degree of influence in the foundation will be based on the degree of involvement in the system, e.g. the amount of AVW token represented by each entity.

The Social Impact

Even though propaganda and information warfare is likely as old as spoken language, the recent rise of social media and echo-chambers enabled by their use has made it much easier to influence public opinion outside of the traditional media. In countries that have a relatively free press, the media often acts as a filter to fact-check outrageous claims and hoaxes propagated by public figures. But the decline in digital ad revenue, combined with the erosion of trust in traditional institutions in a number of countries in the western hemisphere has made it harder to perform this critical role. As a

result of this and other factors, during the last eight years democracy and freedom has seen a continuous decline in U.S and across the word (Fig 7).

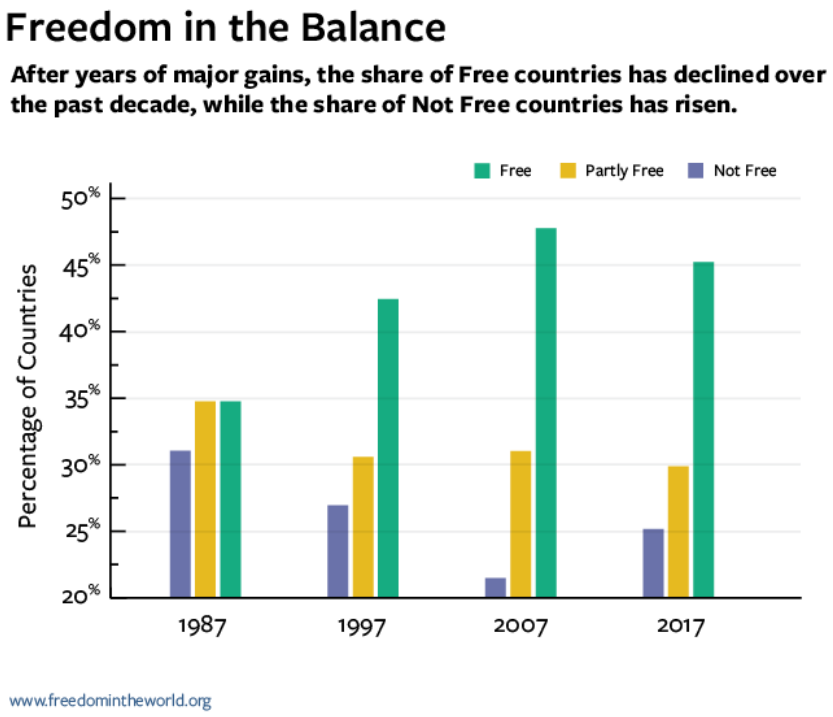

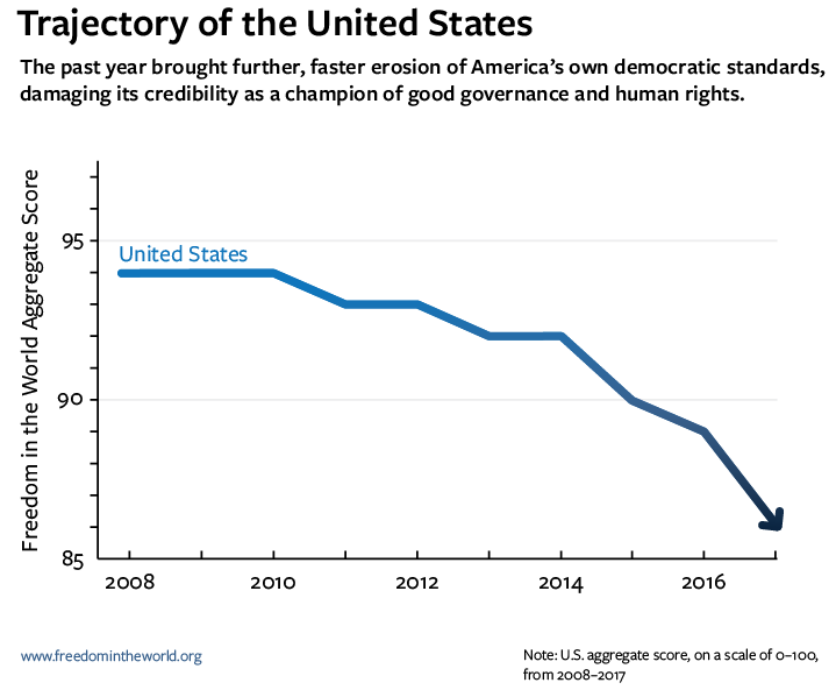

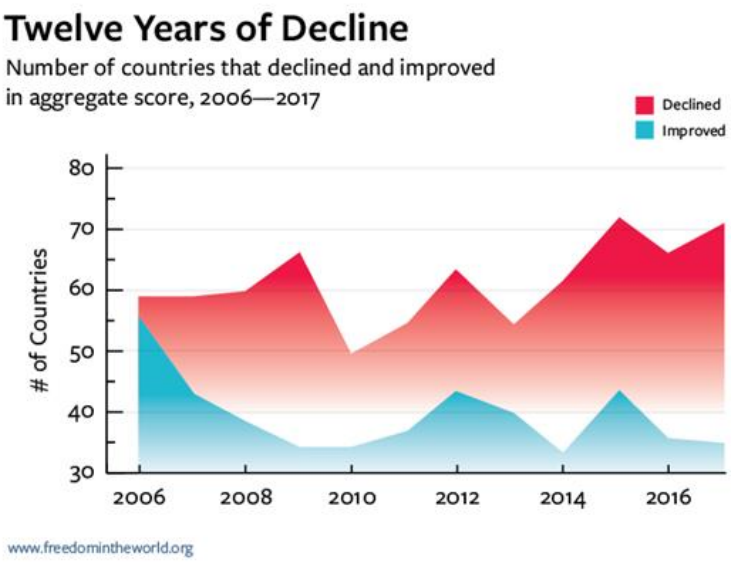

Figure 7. The decline in democracy in US and across the world.

Avow will combat this trend by providing a unbiased, expertise-driven platform for the society to agree upon a set of facts. Such agreed-upon facts are a crucial ingredient in civic discourse. The decentralized nature of the platform will allow the creation of Gateways in different language and with region-specific customization.