

Economic and Business Dimensions Proposal: A Market for Truth to Address False Ads on Social Media

Guaranteeing truth in advertising.

WHEN IT COMES to political ads on Facebook, anything goes. On Twitter, nothing does. In a speech at Georgetown University last October, Mark Zuckerberg, Facebook’s chief executive, defended the company’s decision not to fact-check political ads on the site. Shortly after, Jack Dorsey, Twitter’s chief executive, tweeted that his company had decided to reject all political advertising. In January of this year, Facebook doubled down on its original decision to accept all political ads no matter how egregious the lies an ad buyer wishes us to believe.

They are both wrong. Facebook pollutes our political discourse. Twitter impoverishes it. Between promoting false ads and rejecting truthful ones, here’s a better way: create a “market for truth.” It requires neither machine algorithms to discern truth nor judgments by a potentially self-interested company. Instead, it discourages liars from lying.

First, ask political advertisers to guarantee their truth. Each politician or PAC that places an ad would put a large sum of money in escrow as an “honest ad pledge” that their claims are true. Second, if anyone disputes the ad, an independent fact-checker would judge the ad’s truthfulness. This role could



fall to any one of a number of organizations that routinely make such judgments: FactCheck.org, Politifact, HoaxSlayer, or Snopes. It could even be a panel sampled randomly from Fox and CNN viewers. The watchword is independence. It cannot be Facebook’s self-appointed Oversight Board and it most emphatically cannot be government.

To dispute an ad, an aggrieved party must issue a challenge by paying a non-refundable fee. This challenge price should cover the cost of independent

fact checking. Unchallenged ads wouldn’t require fact checks. If an ad is challenged, the independent fact checker then judges the disputed ad’s accuracy.

If the ad proves false, the injured party receives the advertiser’s pledge, which they can spend as they wish to undo the damage the false ad has caused. If the ad proves true, however, the pledge reverts to the ad buyer, and the challenger forfeits the cost of the fact check. Challengers would have no incentive to pay for fact checking that

Results of ACM's 2020 General Election

President:

Gabriele Kotsis
(July 1, 2020–June 30, 2022)

Vice President:

Joan Feigenbaum
(July 1, 2020–June 30, 2022)

Secretary/Treasurer:

Elisa Bertino
(July 1, 2020–June 30, 2022)

Members at Large:

Nancy M. Amato
(July 1, 2020–June 30, 2024)

Tom Crick
(July 1, 2020–June 30, 2024)

Susan Dumais
(July 1, 2020–June 30, 2024)

Mehran Sahami
(July 1, 2020–June 30, 2024)

Alejandro Saucedo
(July 1, 2020–June 30, 2024)

proves their opponents are right. If an ad goes unchallenged after, say, 30 days, the honest ad pledge reverts to the ad buyer.

In all cases, the cost of guaranteeing the truth of an honest ad is zero. The false advertiser, however, has paid for the ad, paid the pledge penalty, and paid in reputation. Simply put, the forfeited pledge is the price of a lie. It is paid only by liars. A politician who still wishes to lie may certainly do so. But lying becomes expensive.

What about the slippery middle ground between truth and falsehood—the innuendo and half-truths that infect so much political advertising? Imagine a photo of Joe Biden and his son looking shifty, accompanied by the tagline: “Hunter Biden served on the board of Ukraine’s most corrupt company while his father, as Vice President, did all he could to fire a powerful Ukrainian prosecutor.” None of that is exactly false. But it implies the senior Biden tried to prevent the prosecutor from going after the company, when in fact he sought the opposite: he wanted the prosecutor fired for failing to pursue corruption.

How should an honest ads market handle an ad like this? It refunds half the pledge for an ad that’s half a truth. Based on the egregiousness of the lie, the amount of a refund can correspond to one of the sliding scales fact checkers already use. Indeed, Politifact did rate an ad like the one here as half-true on a scale that ranges from: true, mostly-true, half-true, mostly-false, false, and pants-on-fire. Other fact checkers use similar scales. A market for truth need not be perfect. It just needs to be credible and unbiased. By asking PACs and politicians to warrant their claims, it changes the balance of power, favoring truth over lies in our political discourse.

**A market for truth
need not be perfect.
It just needs
to be credible
and unbiased.**

A politician of greater integrity could make bigger promises and voluntarily escrow 2X or 3X the normal pledge. Or, a politician of little means could pre-check a message with the fact checker, ensure an honest message and have bonds markets underwrite the pledge without risk. By contrast, a politician of low integrity could afford only smaller promises, namely the minimum lie price that the media platform requires as its honest ad pledge. And, low integrity politicians would keep losing their pledges. For a dishonest politician, the costs mount with each additional lie.

How might this work? To operate a market for truth, we can rely on established administrative practices that we already use for trust and legitimacy. Taking our own government as precedent, we split oversight into legislative, judicial, and executive branches. A legislative body gets to define “fake ads.” Despite their differences, even Fox News, CNN, and the *New York Times* might be able to agree on a working definition of fake ads independent of specific use cases and their own ads and news stories. A judicial body gets to decide whether a specific case represents an instance of fake advertising according to this definition. Again, Snopes, Hoax Slayer, Politifact, or a jury of peers might play this role only now they must judge according to the definition provided by the legislative body. Jurors do not get to use their own individual definitions. Finally, the executive branch enforces these definitions and decisions. It collects the honest ad pledges and disburses them to ad buyers or ad challengers based on rulings by the judicial body. Social media platforms like Facebook and Twitter can play this role but they decide neither the definitions nor the outcomes of challenges. By dividing the branches of fake ads governance, we recreate an institution where no branch judges truth as applied to itself and no branch has an economic incentive to bias its behavior to get rich.

Why does this work? A truth market for trading honest ads works for exactly the same reason as a carbon market based on cap and trade. It solves the problem of pricing externalities and markets for trade in externalities already exist. Carbon dioxide is pollution. It is a negative externality that harms others. An entity that is causing damage needs

to pay for that damage by buying pollution credits that put a price on the harm done. Fake news is pollution. It is a negative externality that harms others. The size of the honest ads pledge, that is, the lie price, could be any escrow amount set by the social media platform but really should be the expected size of the harm done. This negative externality is the “social cost” of the damage done by lying. The crowdsource identification of harm is the market that “trades” the externality. The harmed parties claim the lie price and get paid for the damage they experience. Carbon trading markets work so we can expect markets for truth will also work.

Importantly, a market for truth works even when the amount of damage, the lie price, is not known in advance. Imagine Exxon Mobile today taking out an ad that human activity does not cause global warming. The lie price for political ads in the U.S. alone is too small for the lie price of global warming policy ads internationally. You can quickly see that, if a firm repeatedly pays the lie price, then their willingness to keep lying is too small relative to the true social cost. Then the lie price should rise until they stop the lies that harm people. In other words, we have an “efficient search” process that can force firms and super PACs to internalize the true social cost of their negative externalities even when that cost is initially unknown.

And what about free speech? In the U.S., skeptics might object that an honest ads pledge would not withstand First Amendment scrutiny if the pledge were mandatory. U.S. courts view impediments to speech as violations of free speech. Although this is a uniquely U.S. problem, the system still works even when a pledge is voluntary. If the market for truth is fully functioning, then unwillingness to pledge an honest ad is itself a signal that the author is likely lying because honest ads incur no added cost. The 2001 Nobel Prize in Economics acknowledged the tenets of information economics precisely due to the power of “signals” to separate truth from lies. Informative signals are potentially expensive actions taken by knowledgeable parties that back up their claims. A product sold with a guarantee, for example, is almost always more reliable than a product sold “as is” or “buyer

Good sellers, knowing their claims are true, can offer guarantees that bad sellers, knowing their claims are false, cannot afford to offer.

beware.” Good sellers, knowing their claims are true, can offer guarantees that bad sellers, knowing their claims are false, cannot afford to offer. The voluntary signal separates good from bad, and fact from fiction. The proposed mechanism is very powerful.

An honest ad pledge discourages political advertisers from placing false ads. The pledge need not be mandatory—advertisers’ failure to pledge signals they do not believe their own claims. A fair challenge price discourages political adversaries from launching false challenges. The mechanism also provides revenue to pay for independent fact checking via issued challenges. Fact checkers have no financial incentive to bias their decisions and not every ad would need checking—only those that are challenged. Governance can proceed using models we already use in other contexts.

As with so many other aspects of social media, platforms like Facebook and Twitter have made the spread of false content what the tech world proudly calls “frictionless.” It is time to judiciously put some friction back. A truth market would do just that. A society that values unfettered expression over truth can set the price of lying low. A society that values greater integrity can set the price of lying higher. Currently, the price of lying in political ads is zero. □

Marshall Van Alstyne (mva@bu.edu) is a Questrom Chair Professor at Boston University where he teaches information economics. He is also a Digital Fellow at the MIT Initiative on the Digital Economy and co-author of the international best-seller *Platform Revolution* (W.W. Norton).

Copyright held by author.

Calendar of Events

At press time, scheduled conferences were significantly impacted by COVID-19, often requiring cancellation or postponement. The following conferences are scheduled to be held virtually. Please check conference websites for updates on programs and event dates.

July 6–10
DIS '20: Designing Interactive Systems Conference 2020,
Sponsored: ACM/SIG,
Contact: Kristina Anderson,
Email: h.k.g.andersen@tue.nl

July 6–10
WebSci '20: 12th ACM Conference on Web Science,
Sponsored: ACM/SIG,
Contact: Dame Wendy Hall,
Email: wh@ecs.soton.ac.uk

July 8–12
GECCO '20: Genetic and Evolutionary Computation Conference,
Sponsored: ACM/SIG,
Contact: Carlos A. Coello,
Email: ccoello@cs.cinvestav.mx

July 8–12
LICS '20: 35th Annual ACM/IEEE Symposium on Logic in Computer Science,
Contact: Lijun Zhang,
Email: zhanglj@ios.ac.cn

July 13–17
ACM EC '20: ACM Conference on Economics and Computation,
Budapest, Hungary,
Sponsored: ACM/SIG,
Contact: Peter Biro,
Email: peter.biro@krtk.mta.hu

July 13–17
DEBS '20: The 14th ACM International Conference on Distributed and Event-based System,
Sponsored: ACM/SIG,
Contact: Kaiwen Zhang,
Email: kaiwen.zhang@etsmtl.ca

July 13–17
PEARC '20: Practice and Experience in Advanced Research Computing,
Sponsored: ACM/SIG,
Contact: Gwen Jacobs,
Email: gwenj@hawaii.edu