

BOTTLENECKS FOR EVIDENCE ADOPTION

Stefano DellaVigna, Woojin Kim,
Elizabeth Linos, and Jeremy Margolis
April 2024



The People Lab



HARVARD Kennedy School

Understanding the barriers to evidence adoption—and how to reduce them—is critical to ensuring that evidence improves policy outcomes. We followed up with 67 U.S. city departments that collectively ran 73 randomized controlled trials (RCTs) from 2015–2019 to answer two main questions: (1) How often are tested innovations actually adopted? (2) What factors predict whether innovations are adopted by governments? Additionally, we conducted a prediction survey among researchers and practitioners to examine whether experts accurately predict the factors that influence evidence adoption.

CONTEXT

The last decades have seen a dramatic increase in rigorous evaluation in government, partly in response to a broadening interest in and desire for evidence-based policy making. As a result, we now have a growing set of RCTs that have provided evidence on “what works” (and what does not) across a wide range of critical policy issues. This evidence-based movement has the potential to dramatically improve policy outcomes over time, but only if policymakers are aware of best practices and choose to adopt them. To date, there has been limited research on how often tested interventions are adopted at scale, nor on what motivates experimentation and adoption in the public sector.

RESEARCH

From 2015 to 2019, the Behavioral Insights Team - North America (BIT-NA) collaborated with cities across the U.S. to conduct RCTs. Of all of the studies BIT-NA conducted during this time period, we identified a sample of 73 RCTs across 67 city departments in 30 cities to include in this study. These RCTs are broadly comparable: the trials all included a clear control group and a binary outcome related to take-up, and tested a light-touch “nudge,” such as developing or revising government communications (e.g., letters or emails).

In 2021, we contacted each city department involved in one of these trials to ask about the adoption of the tested innovation, as well as to gather additional information about the implementation. For all 73 RCTs, we were able to assess whether the innovation was adopted and measure other potential predictors of adoption.

KEY TAKEAWAYS

- 1 Across 73 RCTs conducted in 67 U.S. city departments, we find that just 27% of tested interventions were subsequently adopted.
- 2 Neither state capacity nor the effectiveness of the intervention significantly predicts adoption.
- 3 The strongest predictor of adoption is whether the tested intervention involved modifying a pre-existing communication or developing a new one.

RESEARCH (cont.)

We define successful adoption as any case where at least one innovation or treatment that was tested in the original RCT has since been used in communications from the same city department. We also examine three factors that could influence adoption: (1) the strength of the evidence in terms of both effect size and statistical significance; (2) organizational features such as the state capacity and whether the city staff member who worked on the RCT is still employed at the city; and (3) the type of innovation.

Additionally, we conducted a prediction survey among relevant experts via the Social Science Prediction Platform (N = 118) that we use to contextualize our findings.

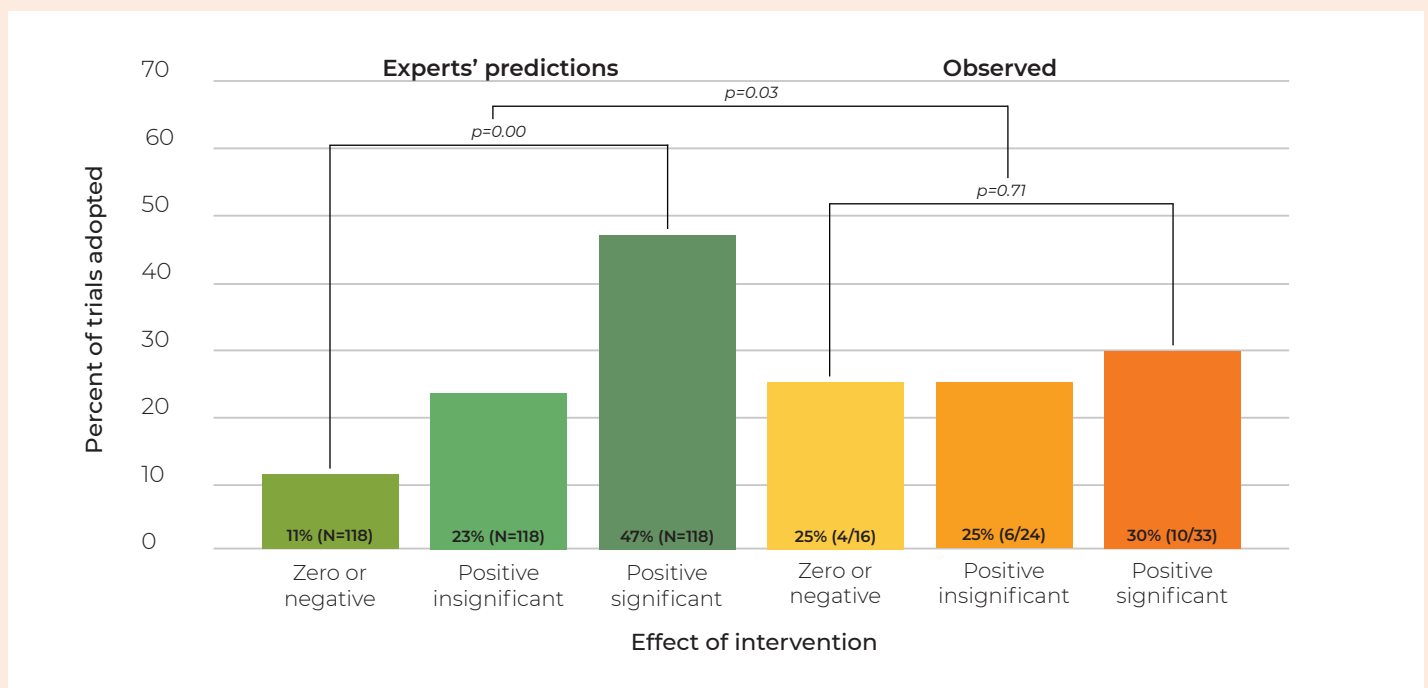
WHAT WE FOUND

Overall, out of 73 trials, we find that the tested innovation was adopted by the city 27% of the time—this is comparable to what experts predicted.

Next, we examine the extent to which different factors predict adoption.

- 1 Strength of the evidence:** Contrary to experts' predictions, we find that the strength of evidence (in terms of effect size and statistical significance) does not significantly predict adoption. Across the 73 trials, interventions that had a negative effect were adopted 25% of the time; those with a positive but not statistically significant effect were also adopted 25% of the time; and those with a positive, statistically significant effect were adopted 30% of the time. These differences in adoption rates are not statistically significant (see Figure 1). We also find that the magnitude of the effect size does not predict adoption.

FIGURE 1
Adoption rates by effectiveness of the tested intervention



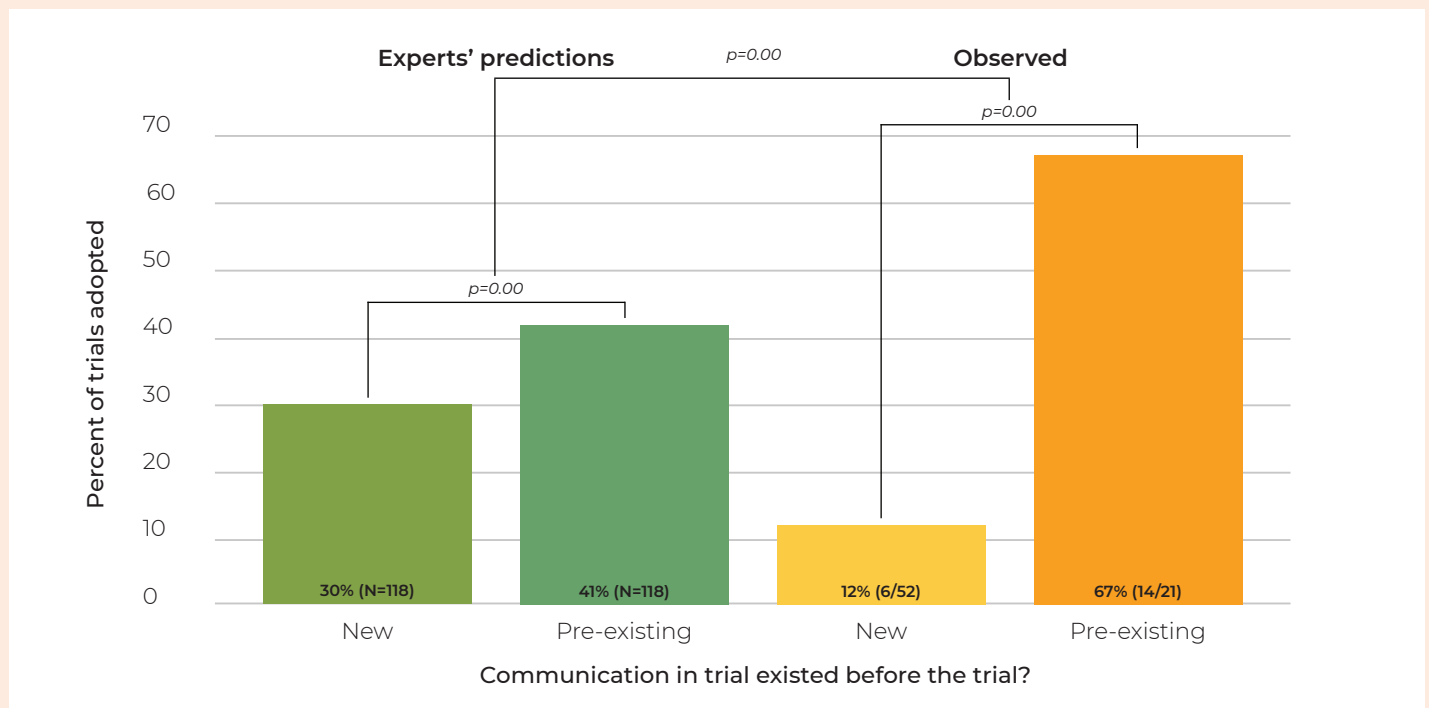
WHAT WE FOUND (cont.)

- 2 Features of the city:** First, we use city population and whether the city had been certified as a data-driven city by the What Works Cities program as proxies for city capacity, but find that neither measure predicts adoption. We also examine whether cities that still employed the primary city staff member who collaborated with BIT-NA on the original RCT were more likely to adopt the intervention. We find that these cities were 14 percentage points more likely to adopt the tested intervention (19% vs. 33%), but this difference is not statistically significant.
- 3 Experimental design:** We find that—by far—the strongest predictor of adoption is whether the tested intervention involved modifying a pre-existing communication or developing a new one. Of trials that involved modifying a pre-existing city communication, 67% were adopted, compared to only 12% of those that involved developing a new communication (see Figure 2). This difference is highly significant and holds under a series of robustness checks, and is also much larger than what experts predicted.

Meanwhile, we find that the choice of behavioral mechanisms used in the intervention (e.g., simplifying language vs. using social cues) has little impact on adoption.

FIGURE 2

Adoption rates by whether the tested intervention modified a pre-existing communication or was a new communication



Why are interventions that modified pre-existing communications so much more likely to be adopted? We identify organizational inertia as a leading explanation: changes to pre-existing infrastructure are more naturally folded into subsequent processes.

WHAT'S NEXT

In exploring the long-term outcomes of 73 RCTs conducted across 67 U.S. city departments, we show that there are substantial bottlenecks to adoption of evidence. We find that the strength of the evidence and key city features do not strongly predict adoption; instead, the largest predictor is whether the RCT was implemented using pre-existing communication, as opposed to new communication. We also show that these bottlenecks to adoption have a meaningful real-world impact: the average impact on policy outcomes of the adopted innovations in our sample was just one-third of what the impact would have been had all of the effective nudges been adopted. These findings suggest that targeting and reducing barriers to adoption could yield higher adoption rates, with a meaningful impact on policy outcomes. Future research should consider and test methods of mitigating adoption bottlenecks.



The People Lab



HARVARD Kennedy School

About The People Lab

The People Lab aims to empower the public sector by producing cutting-edge research on the people of government and the communities they serve. Using evidence from public management and insights from behavioral science, we study, design, and test strategies for solving urgent public sector challenges in three core areas: strengthening the government workforce; improving resident-government interactions; and reimagining the production and use of evidence.



Contact Us

✉ peoplelab@hks.harvard.edu

🐦 [@HKS_PeopleLab](https://twitter.com/HKS_PeopleLab)