Lead Poisoning in Low- and Middle-Income Countries: Causes, Effects, and Solutions

Mikey Jarrell

December 2022

Abstract

Lead is an extremely harmful pollutant, especially to children, having been linked to violence, low IQ, heart disease, and other adverse outcomes — even at low levels of exposure. In rich countries, lead exposure has been greatly reduced, thanks at least in part to government regulation: most notably, bans on lead paint and leaded gasoline. However, despite similar regulations, lead levels are still alarmingly high in poor countries. I ask three important questions about lead in poor countries: (1) Are the existing regulations effective? (2) What are the major sources of exposure? (3) When regulation for a major source is impractical, can we shift its market to a new, lead-free equilibrium? To answer these questions, I propose three studies: (1) An event study using exogenous variation in the timing of leaded gasoline bans across Africa; (2) a randomized experiment that staggers the rollout of a subsidy for lead-free paint; (3) an experiment that cross randomizes publicly available lead-testing equipment with the introduction of a reputable lead-free turmeric vendor. These three studies will allow us to better understand the sources of exposure to lead in poor countries, methods to eliminate that exposure, and the gains from doing so.

Contents

1 Introduction

Lead has been extremely useful to humankind over the millennia. Its high malleability and durability, combined with its low melting point, make it ideal for uses including bullets, pipes, stained glass, ceramic glaze, and even preservatives in wine and liquor (Jashemski and Meyer, 2002; Eschnauer and Stoeppler, 1992). Its widespread use in plumbing is reflected in the word itself: the etymology of the English word "plumbing" traces back to the Latin *plumbum*, meaning "lead" (Abbott, 1877). Unfortunately, lead's usefulness is surpassed only by its toxicity, a fact of which humankind has been aware almost as long as they have been of lead itself. Writers as far back as Hippocrates (460–377 BC) and Vitruvius (80–15 BC) were concerned about lead poisoning and accurately described its symptoms (Lessler, 2006; Hodge, 1981). In high doses, lead's acute harms are well described by the medical case study literature[1].

While acute exposure in the modern era is exceedingly rare, mild lead exposure is ubiquitous. At low levels, effects of lead exposure are "subclinical" (i.e., without obvious symptoms), making it difficult to study: exposure to lead is correlated with innumerable other confounders (e.g., income, age, place of birth, etc.) that are likely to be correlated with any hypothesized outcome. Studies from neuroscience that describe a pathway from lead exposure to adverse outcomes such as impaired impulse control via abnormal prefrontal cortex development are enticing (Stiles and Bellinger, 1993), but these correlational studies fail to rule out reverse causality — what if the poor impulse control caused exposure to lead? — or confounding variables — what if living in poor neighborhoods causes both exposed to more lead (via old housing stock) and poor impulse control (via absense of good role models). Furthermore, there is no experimental evidence on the effect of lead poisoning on humans (for good reason, of course). What, then, allows the World Health Organization to declare that "there is no level of exposure to lead that is known to be without harmful effects" (WHO, 2022)?

[1]See Appendix A.1 for an overview of this literature.

Recent attempts to quantify a causal effect of mild chronic lead poisoning on long-run outcomes rely on highly-controlled cross-sectional correlations, longitudinal studies, and quasi-experimental evidence. Such evidence has (more or less convincingly) identified mild lead poisoning as the cause of crime, behavior problems, low test scores, low IQ, low fertility, low birth weight, and increased mortality from cardiovascular problems.[2] Much of the literature points to high long-run benefits from preventing exposure to lead in childhood, a result that mostly exploits enormous variation over time in lead exposure in high-income countries (HICs).

In HICs, lead poisoning climbed throughout the 20$^{\text{th}}$ century until its peak the 1970s (Nevin, 2000). Around that time, clinical evidence demonstrating the harms of lead exposure had begun to pile up, and public sentiment around lead began to change (Stasik et al., 1969). As a result, HICs began making regulatory changes to counteract lead poisoning. The most important was the phase out of leaded gasoline, but regulation was also imposed around the same time on most other known sources of lead: for example, regulation set a limit on lead in paint, toys and furniture, soldered tin cans, pesticides, and drinking water (Nevin, 2007). The cumulative effect of those regulatory changes was a success: in the US, the share of children aged 1–5 whose blood contained more than 10 micrograms of lead per deciliter of blood (µg/dL) — a threshold that today triggers a intense intervention in most states — fell from 88% in the late 1970s to 0.2% today (Billings and Schnepel, 2018; Dignam et al., 2019).

The situation in LMICs looks a bit like it did in HICs forty years ago. While the data to make such sweeping assertions is disappointingly sparse — only two LMICs have anything resembling a nationally representative sample of blood–lead results, and the vast majority of LMICs have no blood–lead data whatsoever — it is estimated that at least one in three children in LMICs have blood–lead levels above 5 µg/dL, a threshold long referred to by the CDC as "concerning" (Ericson et al., 2021).[3] In total, 99% of children with lead poisoning live in LMICs. If the effect of successfully reducing lead exposure in LMICs is proportional to the effect of having done so in HICs, then it is estimated that $1 trillion per year in social costs can be averted, primarily through boosting IQ (Attina and Trasande, 2013).

The obvious solution, therefore, would be for LMICs to follow the same playbook

[2] See Appendix A.2 for an overview of this literature.

[3] For convenience, when I say "lead poisoning," I am referring to anyone whose blood–lead levels exceed this threshold, but such a distinction is fairly arbitrary, especially considering the WHO's assertion that there is no safe level of lead exposure. On the other hand, most humans alive today have non-zero blood–lead levels, so it is helpful to draw the line somewhere.

as HICs. However, this does not appear to be working. For example, from 2000–06, every country in Africa (except Algeria) banned leaded gasoline. If leaded gasoline was responsible for the rise and fall of lead poisoning in HICs, why are blood–lead levels still so high LMICs? Possible explanations fall into two main categories. The first category asserts that leaded gasoline was not a major contributor to lead poisoning in LMICs — because, for example, there is less vehicle traffic, or people live farther away from roads. The second category asserts that leaded gasoline was a major contributor, but its effects are difficult to detect — because, for example, the effect of the ban is slow to materialize in the data, or because its effect was offset by the rise in other source of lead exposure. I propose an event study to attempt to litigate between these two types of explanations. See Section 2.

Even if leaded gasoline is, in fact, a major contributor to lead poisoning in LMICs, it is unlikely to be the only one. Another problem with the HIC lead-elimination playbook is the lack of political impetus to enforce regulations that are already on the books. For example, lead paint is an obvious potential source of exposure, but despite the existence of lead paint regulations in most LMICs, lead paint is still widely available (Coulter, 2022). It appears that consumers (and the government officials charged with protecting the best interests of these consumers) are unaware that the paint available for sale at the local hardware store is likely well beyond the maximum allowable lead concentration. One US-based non-governmental organization has found success nudging LMIC governments in the right direction: informing them of the prevalence of lead paint, facilitating effective testing and enforcement. They also working with vendors to help them source from lead-free suppliers. How much good this will do depends on how much lead poisoning in LMICs is attributable to lead paint. Is enforcing these regulations an efficient use of severely constrained government resources? I propose a randomized controlled trial that staggers the rollout of a lead-free paint subsidy program, creating a treatment group that is exposed to the effect of the regulation enforcement earlier than the control group, resulting in exogenous variation in exposure to lead paint and, if there is a reasonably strong first stage, exogenous variation in blood–lead levels. To my knowledge, this would be the first experimentally created variation in lead exposure, which will allow us to precisely estimate the short- and long-run causal effect of lead poisoning. See Section 3.

A final problem with the HIC playbook is that there exist some sources of lead exposure that are difficult to regulate. A market for one such source has been studied extensively by Jenna Forsyth (Forsyth et al., 2018, 2019a,b; Forsyth, 2022). In Bangladesh, turmeric is often adulterated with yellow dye to cut costs and improve perceived quality. The dye used is lead chromate, and evidence suggests that this is responsible for the majority of lead poisoning in Bangladesh. Regulation is difficult because the industry is composed of many small players distributed throughout the country, and because testing requires expensive equipment. Plus, there is no political motivation to regulate: consumers only care about turmeric color and appear to be unaware of the health consequences. However, simply shifting consumer preferences appears to be necessary to escape this harmful equilibrium, but it may not be sufficient; it is impossible to distinguish clean from adulterated turmeric with the naked eye, ear, mouth, or nose. Therefore, I propose a two-pronged intervention. The first treatment is to provide access to lead testing equipment, randomized at the market level, that will allow buyers easily test for adulteration. Making quality observable may shift the market to a new equilibrium via a demand shock. On the other hand, a supply shock may also be necessary, so I plan to cross-randomize the lead testing treatment with a reputable seller treatment: the entry of a new turmeric brand that makes a credible claim guaranteeing no adulteration. It is plausible that the two treatments are complementary: the testing may only effective in a market where sellers advertise being lead-free, and the advertising may only be effective if the buyer can verify sellers' claims. See Section 4.

In light of the enormous potential gains from reducing lead poisoning in LMICs, it is essential that we better understand what are the primary sources of lead poisoning in LMICs, which interventions can curtail those sources, and whether those interventions improve long-run outcomes. The rest of the paper is organized as follows. Section 2 describes my proposal to conduct an event study of leaded gasoline bans in Africa, Section 3 describes my proposal to conduct a randomized evaluation of a lead paint mediation program, Section 4 describes my proposal to conduct a randomized evaluation of introduction of an honest turmeric seller along with capacity to test turmeric for lead in Bangladesh, and Section 5 concludes. Appendix A conducts a full review of relevant literature, Appendix B contains information relevant but not essential to the proposals, and Appendix C describes other promising avenues for research into the causes and effects of and solutions to lead poisoning in LMICs.

2 Proposal I: Event study of ban on leaded gasoline in Africa

There is a growing literature on the effects of lead on crime (see section A.2). Most studies on the lead–crime effect, such as Reyes (2007), use variation caused by the phase-out of leaded gasoline, which appears to explain a large fraction of the falling rates of violent crime in HICs across the past few decades since its peak around 1990. Whether this effect exists is well established, but whether its magnitude is meaningfully large is still hotly contested (Higney et al., 2022). Competing hypotheses for explaining the falling violent crime rates in HICs include falling poverty, demographic transition, better policing, increased access to abortion, and many others (Rosenfeld and Fornango, 2007; Baumer et al., 2012; Levitt, 2004; Donohue and Levitt, 2019). Tcherni-Buzzeo (2019) summarize 24 of these different explanations proposed in the literature, and conclude that the lead–crime link is a promising hypothesis that requires further research.

Furthermore, as far as I'm aware, none of the studies on this topic come from LMICs. While it is reasonable to expect that the effects of lead on crime would be similar across countries, it is not clear that the effect of banning the sale of leaded gasoline would be similar across countries. In fact, available data, sparse though it may be, suggest that there are still remarkably high blood–lead levels present in LMICs today, so it is plausible that banning leaded gasoline has had little effect in those countries (Ericson et al., 2021). Might it be the case that people in LMICs are exposed to less vehicle exhaust than people in HICs? On the other hand, because there is necessarily some delay between the banning and leaded gasoline and the appearance of downstream effects, and because leaded gasoline was banned much more recently in LMICs than in HIC, it could also be the case that insufficient time has passed between the ban and the collection of blood–lead data to which we have access.

Up until 2000, leaded gasoline was in use in every country in Africa. One by one, countries began banning it, mostly in response to pressure from the United Nations (Lacey, 2004). By 2006, every African country had banned leaded gasoline — except Algeria, which didn't act until 2021. It appears that the exact timing of ban from country to country within Africa is random, although some research will be required to test the veracity of this assumption. Over time, the ban prevented the release of lead particulate matter into the air, which, presumably, caused blood–lead levels of people breathing that air to be lower than it would have been absent the ban.

So why the delay in effects? As discussed in Appendix A.1, the harms of lead poisoning are felt much more severely in children than in adults. Because one's propensity to commit crime tends to peak around the ages of 15–24 (Ulmer and Steffensmeier, 2014), we should not expect the ban to have an immediate effect on crime; rather, the effect should grow slowly over time as more of those in their peak crime-committing years have spent more of their childhood in reduced-lead environment. This is not dissimilar to the phenomenon similar found in Duflo (2001), where the new schools were only beneficial to those young enough to attend them, so the effect of the schools on schooling grows as we examine younger and younger cohorts; however, the effects of schools on income doesn't manifest until those children grow up and enter the labor force.

Just as exposure to the effects of the ban varies with time, it also varies with space. People who live near a high volume of vehicle traffic are exposed to a high volume of car exhaust, so the removal of lead from gasoline should have caused a large reduction in lead exposure in these "high-intensity" areas. Conversely, people living areas with little vehicle traffic are much less exposed to car exhaust, so the ban should have little effect on their lead exposure in the "low-intensity" areas. I plan to estimate the effect of banning leaded gasoline on crime in Africa by comparing the change in crime before and after the ban in high-intensity areas to the change in crime in low-intensity areas.

2.1 Data

For our outcome variable, what we need is something measured with sufficient spatial and temporal resolution to be able to detect country-level trends year over year. Afrobarometer provides an ideal data set, featuring geolocated measures of crime and crime-adjacent outcomes such as perceived safety. I will construct a standardized index of perceived crime, using questions such as "Over the past year, how often, if ever, have you or anyone in your family felt unsafe walking in your neighborhood?" and "During the past year, have you or anyone in your family been physically attacked?"

Our measure of exposure to vehicle exhaust by region will be constructed using several possible sources, including nightlights, gasoline sales, road density, population density, and remote sensing data on CO_2, NO_2, and other particulate matter found in car emissions. One promising technique, which uses a random forest classifier on data from the Sentinel 2 satellite, has been used to precisely count the number of trucks on the road in England and in Kenya (Douglas, 2021). A drawback with this approach is that we're interested in exposure to vehicle exhaust around the time of the ban, while this data tells us about much more recent conditions. While it is likely that traffic levels today are highly correlated with traffic levels 20 years ago, it is difficult to rule out trends that deferentially affected regions before and after the ban.[4]

For timings of the ban on leaded gasoline by country, we will rely on research conducted by Our World in Data (Ritchie, 2022). More research will be conducted into how quickly the bans were enforced in each country. For example, was leaded gasoline slowly phased out, as it was in the US, or was the transition much faster as a result of the rest of the world already having transitioned? In either case, it is likely that treatment is not binary; rather, it should increase in intensity over time.

2.2 Estimation

To simplify exposition, let's begin with a standard two-way fixed effects event-study model,

$$Y_{itr} = \alpha_t + \lambda_r + \sum_{\tau=-6,\tau\neq-1}^{22} \beta_\tau \mathbb{1}(t - b_r = \tau) + \epsilon_{itr}, \tag{1}$$

where Y_{itr} is a constructed index of crime as perceived by person i at time t in region r, α_t is survey year fixed effects, λ_r is region-level fixed effects, and $\mathbb{1}(t - b_r = \tau)$ is an indicator for whether year t is τ years after the year of the ban b in region r. The coefficients of interest here are represented by β_τ, which tell us the effect of the ban on crime τ years after the ban. If our hypothesis — that there is a several year delay between the ban and its effect on crime, and the effect increases with time — holds, we should expect β_τ to be indistinguishable from zero for at least $\tau \in \{-6, -5, \ldots, 0\}$. Thereafter, we should expect β_τ to be increasing with τ for $\tau > 0$.

We now can add to the model the constructed measure of exposure in region r to the treatment, which we'll call T_r. For simplicity, for now consider T_r to be a dummy for whether region r is "high intensity," i.e., whether region r has high exposure to car exhaust.[5] In low-intensity regions, set $T_r = 0$.

$$Y_{itr} = \alpha_t + \lambda_r + \sum_{\tau=-6,\tau\neq-1}^{22} \beta_\tau T_r \mathbb{1}(t - b_r = \tau) + \epsilon_{itr}. \tag{2}$$

[4]For example, let's imagine a region home to an informal settlement that experiences rapid increase in population density after the ban. Looking at recent data, this region would be classified as "high-intensity" region. It is easy to imagine this region also experiencing an increase in crime as a result of the increased density; however, my analysis would ascribe this effect to the ban.

[5]When actually estimating the model, depending on data quality and availability, it may be possible to divide T_r into several intensity bins, or even treat it as a continuous measure of intensity.

Interacting our treatment dummy T_r with our event study–time indicator $\mathbb{1}(t-b_r = \tau)$ means that we now have a difference-in-difference-in-differences (DDD), a.k.a., a triple difference design (Muralidharan and Prakash, 2017). This changes the interpretation of our coefficient of interest: now, each β_τ tells us the effect on crime relative to baseline (this is the first difference) in one region compared to regions in different event times (second difference) changes in high-intensity regions compared to low-intensity regions (third difference). In other words, β_τ is the difference between effect of the ban on crime in high-intensity regions and its effect in low-intensity regions. In this model, because intensity is a dummy, low-intensity regions can be thought of as "never treated."

Ceteris paribus, we should see people in high-intensity regions converge to people in low-intensity regions as lead exposure converges towards zero everywhere. This appears to have been the case in the US, where gasoline was far and away the largest contributor to lead poisoning (Higney et al., 2022). However, it is not clear that *ceteris* is *paribus* in this context. High-intensity regions differ from low-intensity regions in many ways other that intensity of car traffic: high-intensity areas will be more densely populated, more urban, higher income, etc. Of course, these differences are controlled for with our region-level fixed effects λ_r; however, if these differences manifest themselves in non-parallel trends, our estimates of β_τ will be biased.

For example, it is known that informal recycling of lead-acid batteries is a significant source of lead exposure in LMICs (Tanaka et al., 2022). This is typically an urban activity (Belay et al., 2015). If the effect of this industry does not vary by year, then it will be controlled for with λ_r. However, it has been found to be increasing in recent years (Tür et al., 2016). If the effect of this increase in lead from this industry is to increase the crime rates of those living near it, then crime in our treatment regions should be trending up with time. This would counteract the effect of the ban of leaded gasoline, leading us to underestimate β_τ. Thankfully, such an effect would also be visible in the pre-trends: as we go back in event time, we would see crime decrease in the treatment regions. We could control for the pre-trend by assuming the effect of battery recycling continues after the ban just as it had before the ban using the extrapolation method described in Rambachan and Roth (2022), but there is no way to test this assumption; if battery recycling suddenly spiked in treatment regions after the ban, then our estimate of β_τ would be irrevocably biased.

The potential for bias is exacerbated by the fact that there is variation in timing of the treatment. As documented by Sun and Abraham (2021), our estimate of β_τ

may be contaminated by cohort-specific effects unless treatment effects are completely homogeneous across cohorts. This is almost certainly not the case in our setting: the effect of the ban on crime depends on the age distribution in the population, and it is unlikely that distributions of ages does not differ across regions. Therefore, we must further interact our model.

$$Y_{itr} = \alpha_t + \lambda_r + \sum_{c=2001}^{2006} \sum_{\tau=-6,\tau\neq-1}^{22} \beta_{\tau c} C_c T_r \mathbb{1}(t - b_r = \tau) + \epsilon_{itr}. \tag{3}$$

Now our coefficient of interest is indexed along two dimensions: event time τ and cohort c. The new dummy, C_c, is equal to one when the cohort C for person i is equal to c, where cohort is defined as the year in which leaded gasoline was banned in region r. There are seven cohorts in our sample — $C \in \{2000, 2001, \ldots, 2006\}$ — but we need to leave one out to avoid collinearity. In order to recover the average treatment effect, we can collapse our various $\beta_{\tau c}$ across cohorts within event time using estimated weights according to share of the sample within each cohort c for each event time τ. Graphing β_τ by event time τ will illustrate the average effect of ban of gasoline on perceived crime over time from several years before the ban until up to 20 years after.

2.3 Alternative strategies

Depending on data availability, a similar approach to that outlined here could be used to estimate the effects of the banning of leaded gasoline in other parts of the world. For example, the ASER Centre has data on test scores in India dating back to 2008, which provides an excellent opportunity to estimate the effects of lead on educational outcomes in a LMIC, which would also be novel.

3 Proposal II: Lead Paint RCT

The Lead Exposure Elimination Project (LEEP) is a recently founded US-based non-governmental organization that has found success removing lead paint from the market in LMICs using a fairly simple strategy.

1. Market analysis: test available paint brands and colors for lead.

2. Government outreach: report results and help develop enforcement strategies.

3. Industry outreach: facilitate production and sale of lead-free paint.

Their approach was first tried in Malawi, and within six months, the sale of lead paint in the country had been stopped. In personal conversations, LEEP founder Lucia Coulter has explained to me that although lead paint is typically a similar price as lead-free paint, it is slightly more durable and vibrant, so ignoring the externalities, the incentives are for the vendors to source from suppliers that use lead[6] (Coulter, 2022). However, LEEP has found that simply testing the paint for lead, showing the results to the vendor, explaining the harms, and connecting the vendor to lead-free suppliers has been sufficient to change behavior. In addition, LEEP has found governments to be responsive to their recommendations and expertise on how to effectively enforce their (already-existing) regulations.

Although it is unequivocally a benefit to health to reduce the availability of lead paint (*ceteris paribus*), the magnitude of this benefit is unknown. There is an appalling dearth of source apportionment studies (studies that estimate shares of lead poisoning attributable to different sources), and as a result, the share of lead poisoning caused by lead paint in any given LMIC could be 0%, 100%, or anywhere in between. Knowing this share would inform policymakers and philanthropic ventures where their limited resources and attention should be directed.

One way to estimate this share is to measure the causal effect of removing lead paint from the home. Absent the possibility of vacating the house so that a professional service can carefully remove the old paint without creating massive amounts of lead-filled dust, the best practice is to leave the old coat undisturbed and paint over it with a fresh coat of lead-free paint (Coulter, 2022). A randomized control trial in which the treatment group is encouraged to do exactly this will allow us to estimate its effect on blood–lead levels, which serves as a convenient and plausible proxy for strict enforcement of a lead paint ban: the control group will feel the effects of the ban once they get around to repainting their home with what will by that point surely be lead-free paint, which should be, on average, well after the treatment group has done so. Taking inspiration from Hamory et al. (2020), an experiment that creates exogenous variation in blood–lead levels presents a perfect opportunity to track participants over time and estimate the long-run causal effects of lead poisoning. This would be a groundbreaking study: to my knowledge, it would be the first that experimentally induces random variation in lead exposure in humans.

[6] Notably, conversations with LEEP have revealed that some of the largest suppliers of lead-based paint pigments come from the UK and Canada.

3.1 Intervention: Subsidized paint, plus a tutorial

We plan to select a country in which LEEP has recently successfully eliminated the sale of lead paint (e.g., Malawi). A baseline survey of a representative sample is required to assess the prevalence of lead paint on the walls of homes, and the blood–lead levels of household members under the age of 12. Households randomized into the control group will hear the results of their blood–lead tests and receive recommendations based on those results, including best practices for safely painting over old lead paint. Households randomized into the treatment group will undergo the exact same procedure as the control group (baseline survey, results provision, painting recommendations), but then they will also receive receive a coupon for a subsidized can of paint at their local shop (conditional on their walls testing positive for lead paint). Control households that test positive for lead paint will also receive the subsidy, but at a later date. Randomization will be clustered at the school level, which is necessary because there may be significant spillovers from treatment children onto their peers via the propensity for children with lead poisoning to be disruptive in the classroom (Gazze et al., 2021).

3.2 Data

Lead paint can easily and cheaply be tested for using 3M LeadCheck Swabs, which give results on the spot. Measuring blood–lead levels is slightly more complicated. A point-of-care device called LeadCare II is a relatively inexpensive blood analyzer that requires a simple prick of the finger. Do to low cost and high convenience, its use is widespread (Rosales-Rimache et al., 2022). However, experts in the field have cautioned me that LeadCare II and other comparable point-of-care diagnostic devices lack the precision that would be required to detect an effect size on the order of what we would expect to find here. Therefore, they have recommended collecting blood samples and having them tested for lead at a lab.

In addition to blood–lead levels and presence of lead paint, baseline surveys will collect data on IQ using Raven's Matrices; on discount rates by adapting for children the method pioneered in Bartos et al. (2021); on executive control using a go/no-go task as is standard in the cognitive psychology literature (Wiebe et al., 2012), and impaired attention span (a proxy for ADHD) using a sustained attention task. In the long run, we will also collect data on employment, income, and health outcomes such as heart disease. Finally, administrative data will be required for school attendance, behavioral problems, and test scores. In addition, in order to collect better data on the take-up of the subsidy (e.g., exact timing), administrative data can also be collected from the paint vendors on coupon use and cans of paint sold.

Because it is uncertain how long the lag is between painting over lead paint and seeing a reduction in blood–lead levels — the delay depends on how long on it takes for the paint to chip and create dust, how long it takes the dust to be ingested, and the bioavailability of the lead in the paint (i.e., how much ingested lead reaches the bloodstream) — followup surveys will have to be conducted at fairly frequent, regular intervals. Most of these can be much more basic than the baseline survey. For example, two months after the baseline survey, a followup survey can be conducted to simply check whether walls have been repainted and to test blood–lead levels of the children.

3.3 Estimation: IV

Consider a simple model of the effect of lead on some outcome — take IQ, for example.

$$Y_i = \alpha + \beta L_i + \epsilon_i \tag{4}$$

Here Y_i is IQ for person i, α is a constant, L_i is the blood–lead level for person i, and β is the effect of blood-lead on IQ. This is perhaps an oversimplifcation of this relationship: there is evidence that the relationship between blood–lead levels and IQ is not linear — to be specific, the dosage response function may be decreasing, i.e., lead has decreasing marginal harms (Lanphear et al., 2005). There are a number of ways of dealing with this, but, for simplicity of exposition, we will do so non-parametrically by allowing for the effect of lead vary flexibly across discrete ranges, which I will call "bins."

$$Y_i = \sum_{k=1}^{\lceil \max(L) \rceil} \beta_k \mathbb{1}(k-1 \leq L_i < k) + \epsilon_i. \tag{5}$$

Here, we divide possible blood–lead levels L into bins indexed by k, where bin k covers the interval $L \in [k-1, k)$, and $\mathbb{1}(k-1 \leq L_i < k)$ is a dummy for whether person i has blood–lead level L_i within bin k (measured in µg/dL). Our coefficients of interest, β_k, index the effects on IQ of having a blood–lead level within bin k compared to omitted group, a blood–lead level in the range 0–1 µg/dl, to which we hereafter refer as the baseline bin.

Of course, estimating this equation directly does not tell us whether blood–lead level has a causal effect on IQ: it is very plausible that a confounding variable (e.g., income) causally affects both IQ and blood–lead levels (omitted variable bias), or that IQ causally affects blood–lead levels (reverse causality). However, because our intervention is a randomized control trial, we experimentally induce exogenous variation in blood–lead levels. Instead, we can estimate a first stage equation that looks like this:

$$L_i = \delta_i + \phi D_i + \nu_i. \tag{6}$$

Here, δ_i indicates individual fixed effects, which allow us to control for baseline blood-lead levels. Therefore, ϕ tells us the effect of treatment assignment D for person i on change in blood–lead level L. Then, we estimate the reduced form,

$$Y_i = \gamma_i + \rho D_i + \mu_i, \tag{7}$$

where γ_i are individual fixed effects, which allows us to control for baseline IQ. Therefore, ρ tells us the effect of treatment assignment D on change in IQ Y for person i. While estimating this equation is useful to policymakers in that it tells us the causal effect of the intervention, what may also be of interest to researchers is the causal effect of lead on outcomes such as IQ. For this, we can estimate a two-stage least squares equation *à la* Angrist (1990):

$$Y_i = \alpha + \beta \widehat{L}_i + \epsilon_i, \tag{8}$$

where $\widehat{L}_i$ is the fitted value of blood–lead level as predicted by baseline blood–lead level δ_i and treatment assignment D_i:

$$\widehat{L}_i = \delta_i + \phi D_i \tag{9}$$

However, as described above, it is necessary to allow β to vary flexibly by bin k, so our two-stage least squares equation that we estimate is the following.

$$Y_i = \sum_{k=1}^{\lceil \max(L) \rceil} \beta_k \mathbb{1}(k-1 \leq \hat{L}_i < k) + \epsilon_i. \tag{10}$$

Here, β_k should tell us the causal effect on IQ of having a blood–lead level between $k-1$ and k. The causal interpretation here is dependent on the exclusion restriction, which stipulates that assignment to treatment affects the outcome only via the instrument. I do not foresee any challenges to the exclusion restriction in this experiment: it seems unlikely that giving somebody free paint should affect their IQ by any pathway other than causing them to paint over lead paint, which would causes the amount of lead in the blood to drop, which would cause their IQ to increases. However, a displacement effect is not out of the question: if the intervention causes demand for paint to increase, raising the price of paint or causing stock outages, the control group would then be subject to a negative spillover effect, upwardly biasing our estimates for the effect of the first stage ϕ. However, this does not affect would be offset by an equally biased estimate of the reduced form ρ, resulting in an unbiased estimate of β_k. In any case, a treatment effect large enough to cause such spillovers strikes me as implausible.

One final note on interpretation is that what we are estimating here is not the treatment effect but the local average treatment effect (Angrist, 1990): the effect on IQ of change in blood–lead levels among people for whom being assigned to the program caused them paint over their lead paint–covered walls. If we have reason to suspect that the relationship between change in blood–lead levels and change in IQ is different for these people than for people who don't bother to paint over their walls, or people who don't have lead paint on their walls to begin with, then that limits the external validity of the findings. However, because this is a purely biological effect, it is unlikely that the IQ effect of lead will be systematically across these different types of treatment response groups.

3.4 Alternative strategies

It is also possible that offering coupons for free paint is insufficient to induce significantly different behavior between the treatment and the control group. Behavioral concerns (e.g., procrastination) make lack of take-up a very real possibility. Gazze (2022) has found a massive effect of ordeals specifically in the context of lead: on average, an extra 15 minutes of commute time (each way) decreased the likelihood of bringing a child in for lead screening by 9%. Piloting the study may resolve the question of take-up, in which case it might be necessary to make a more drastic intervention, such as delivering paint directly to households, or even contracting with painters to repaint the home for free. Otherwise, using nudges (e.g., SMS reminders, disseminating information through 'information hubs,' etc.) may be a more cost-effective method to increase take-up (Banerjee et al., 2021).

If lead paint is used in schools, that might dilute the effect of our intervention, or, at the very least, add noise. Testing schools for lead paint will allow us to determine whether perhaps we should conduct this intervention at the school level instead of or in addition to the household level. If children in this context are spending a large fraction of the day in school, this might make the intervention even more cost effective.

4 Proposal III: Adulterated Turmeric RCT

Turmeric, known for its vibrant yellow color, is a signature spice of South Asian cuisine. The turmeric root, which resembles ginger, is primarily grown in the Indian subcontinent, where it's thought to have originated (Prasad and Aggarwal, 2011). After turmeric farmers harvest the turmeric roots, they must be boiled, dried, polished, and ground into powder. Consumers in Bangladesh say that a bright yellow color is the most important indicator of high-quality turmeric (Forsyth et al., 2019a). In order to increase the perceived quality of their turmeric, vendors often adulterate their roots with a yellow dye. During polishing, the outer skin of the turmeric root is removed in order to expose the bright yellow interior. Adding the dye serves three related purposes:

1. improve the color of the outer layer;
2. reduce the time spent and, therefore, weight lost to polishing;
3. increase the weight (since the dye is cheaper per pound than turmeric).

Unfortunately for the Bangladeshi consumers, the dye that is used is lead chromate, which is intended to be used to color furniture and various plastic products — not food. Although the lead in lead chromate is not thought to be as readily absorbed by the body as other forms of lead, it is present in high enough concentrations in turmeric samples in Bangladesh to be extremely concerning. Recent studies, including one that compares the isotopic signature of lead in blood samples to the signatures of lead in possible sources of exposure in these villages, provide compelling evidence that lead chromate dye is main source of lead poisoning in Bangladesh (Forsyth et al., 2018, 2019b).

There appears to be low willingness of the government to regulate this industry. In contrast to turmeric destined for export, which received the attention of media and government regulators in response to recalls of Bangladeshi turmeric in the US around

2012, industry players report never having heard of even a single case of enforcement of lead regulations on domestic turmeric (Cowell et al., 2017; Forsyth et al., 2019a). The political economic reality is such that the costs clearly outweigh the perceived benefits of regulation enforcement: the prospect of auditing many small players in a broad, diffuse network using extremely expensive equipment in order to achieve vague, diffuse, largely imperceptible benefits many years in the future is doomed by behavioral biases. Testing turmeric for lead has a high fixed cost: conversations with industry experts have revealed that the only viable option for use in the field is a handheld X-ray fluorescent spectrometer ("XRF gun"), which costs around $20,000 (Forsyth, 2022). Fortunately, the marginal cost is very low: it will give an extremely accurate reading of surface-level lead on any object — including turmeric, ground and whole — within a second. But the fixed cost is apparently prohibitively high for the Bangladeshi government.

It is clear why the market supports adulteration: there is no incentive for the vendors to provide a high-quality product, i.e., unadulterated turmeric. In fact, even if it weren't cheaper, consumers would still prefer the adulterated turmeric due to its superior color. In other words, willingness to pay for health is negative in this context. The two different reasons customers don't demand quality imply a need for two different interventions.

1. There is no way for the buyer to ascertain quality. Unlike in Bai (2016), where quality (sweetness) of watermelons is a hidden attribute that is revealed to the consumer after taking a bite, high-quality unadulterated turmeric powder is indistinguishable from adulterated turmeric both before and after purchase: there is no telltale smell, taste, or appearance. This means that the solution proposed by Bai (2016), where the sellers are given access to a costly and difficult-to-forge signal of quality, is less likely to be effective in this context because a consumer has no way of ever finding out whether that signal was accurate. This is an even more extreme scenario than that of Björkman Nyqvist et al. (2022), where consumers have at least some possibility of ascertaining quality. Here, mere entry of a firm selling high-quality goods may not be able to drive the incumbent firms to shift behavior without also making quality observable.

2. Even if buyers were capable of ascertaining quality, it is not clear that consumers would value it. Color is reported to be the most desirable attribute of turmeric (Forsyth et al., 2019a), and behavioral biases may render lead poisoning an

easily dismissed, uncertain future cost compared to the immediate, tangible benefit of the bright yellow color imbued by the dye. This may be an even more extreme scenario than that of Cohen and Dupas (2010), where consumers at least have a non-negative willingness to pay for preventative health. Here, making quality observable may not be able to shift consumer demand without also shifting consumer preferences.

Therefore, I hypothesize that in order to be successful, an intervention here must do two things simultaneously: increase observability of turmeric quality and therefore capacity for reputation building, and increase demand for quality turmeric. If successful, the market will shift to a new equilibrium, allowing us to answer several questions with important policy implications, most notably how to regulate an industry where both supply and demand are working against long-run societal welfare, and where simply imposing a tax to correct for internalities, as studied in Allcott et al. (2019), is infeasible. In addition, this will allow us to test the external validity of the findings of Bai et al. (2022), which found spillovers of a negative shock on other exporting firms — are there also spillovers of positive shocks on domestic firms?

4.1 Interventions

I plan two cross-randomize — following the cross-randomization best practices prescribed by Muralidharan et al. (2019) — two different treatments at the level of markets (a.k.a., "bazaars"): one that allows for buyers to test their their turmeric for lead, and another that introduces a reputable firm selling lead-free turmeric.

4.1.1 Intervention I: Testing for lead with XRF guns

To address the challenge of reputation building for turmeric sellers, I plan to randomize at the market level the provision of an XRF gun at a publicly accessible stand within the market in order to make it convenient for consumers to test turmeric (or other spices) for lead.[7] In theory, absent any other intervention, comparing treatment and control markets should test the hypothesis that consumers already demand quality, and all that was missing was the ability to ascertain it.

[7] In reality, XRF guns are far too expensive to supply to all of the treatment markets in our planned sample. If instead we made XRF testing available at each market only a fraction of the time — say, once per month — then one gun can supply multiple markets with testing capacity at a fraction of the cost.

However, it seems probable that merely setting up a stand in the market brings attention to the fact that turmeric may be contaminated: *if the dye isn't bad, why are those people testing for it?* Therefore, I plan to implement two different treatment groups (still randomized at the market level) in addition to the control group.

1. **Market Testing:** Lead testing equipment at the market (as described above).

2. **Nearby Testing:** Lead testing equipment at nearby government facility.

3. **No Testing:** No lead testing equipment nearby.

By partnering with the government and supplying XRF guns at a nearby government facility (e.g., a police station), we increase an area's capacity for testing without calling so much attention to the fact that testing is happening. This is perhaps more analogous to an under-resourced government attempt to enforce a regulation: a consumer can still have their turmeric tested for lead tested if they really want to, provided they're aware of the service, but there are hoops through which to jump, making the threat of enforcement more hollow than desirable. This allows us to separate out the capacity effect from the salience/convenience effect.

4.1.2 Introducing high-quality seller

The Testing treatment will be cross-randomized with a treatment where we market a new brand of turmeric. By partnering with BRAC, the Bangladesh-based NGO, I plan to bring to market a "lead-free" brand of turmeric. BRAC will contract directly with millers to supply ground turmeric, stipulating that the turmeric must be free of lead. BRAC will have access to XRF guns, allowing them to verify that their suppliers are, in fact, obeying that stipulation. The turmeric will be sold in four different types of packaging, randomized at the market level.

1. **Brand:** BRAC-branded packaging.

2. **Brand Claim:** BRAC-branded packaging that includes a salient lead-free claim.

3. **Warranty:** BRAC-branded packaging that includes a salient lead-free claim and a better-than-money-back guarantee in case of a positive lead test.

4. **Generic Claim:** Unbranded packaging that includes a salient lead-free claim.

The Brand Claim treatment arm packages uses identical branding to the Brand treatment arm, except the Brand Claim packaging features a statement prominently displayed that says something to the effect of, *"Many sellers add yellow dye to their turmeric, which is extremely toxic, and can even cause permanent brain damage in children. We get our turmeric directly from the farmer, ensuring that it is 100% free of lead."* Comparing the Brand treatment arm to the Brand Claim treatment arm tells us how making a claim about the product being lead-free effects demand. There are two possible interpretations of that effect, depending on how the effect interacts with the cross randomized Testing treatments. If it is the case that simply increasing awareness about turmeric adulteration is sufficient for increasing demand for lead-free products, and the BRAC name is endowed with enough positive reputation to give that claim some bite, then we should find more demand for Brand Claim than for Brand, and there should be no interaction with Testing. However, if it is the case that the claim is valueless without some sort of verification mechanism, then we will find demand for Brand Claim exceeding demand for Brand only in markets that receive Market Testing or Nearby Testing — presumably more so in the former than the latter.

The packaging used in the Generic Claim treatment arm is virtually identical to that used in the Brand Claim arm, except nowhere does it say "BRAC" on the generic packaging. Comparing Generic Claim to Brand Claim answers the question of whether the strong reputation of BRAC is necessary to induce a demand response. If the demand is indistinguishable across these two arms, then we can conclude that simply making the claim is sufficient. A plausible result would be that Generic Claim and Brand Claim are indistinguishable only in markets that receive the Testing treatment, where it is possible for any brand to build reputation, whereas in No Testing markets, only a brand that has a pre-built reputation can take advantage of a claim.

The Warranty arm uses packaging identical to the Claim packaging, but it adds a statement to the effect of, *"If turmeric in this package tests positive for lead from a certified facility, you can return it to the seller along with proof of the test results and the seller will give you your money back five-fold."* Such a statement creates an incentive to test the turmeric for lead, increasing the probability of detecting a cheater. Simply adding the threat of (self-)punishment for lying could increase the credibility of the claim, increasing demand for Warranty compared to Brand Claim. Again, it seems likely that the Warranty effect might interact with Testing: if the consumers cannot test the produce for lead, a warranty is nothing more than an empty promise.

4.2 General equilibrium effects

Up until this point, we have only in demand for the products offered in the various treatment arms. These are the direct effects: how do consumers react to the intervention? Perhaps more importantly are the indirect effects: how do other sellers react to reactions to the intervention?[8] The ultimate goal, of course, is to eliminate adulteration from the market entirely, which entails forcing adulterating firms to stop adulterating, or driving them out of business. In order to how the different treatment arms might bring us close to that goal, will use the canonical collective reputation model in Tirole (1996). For a detailed exposition of that model, and what it predicts about the effects of the different treatment arms, see Appendix B.

4.3 Data

There are three categories of data to collect here: data on buyers, data on our new entrant seller, and data on incumbent sellers. For the buyers, we will need to conduct surveys to collect data on turmeric consumption, and the decision-making process behind its purchase. In addition, because we believe the treatment will have a causal effect on lead exposure, it would be nice to collect data on lead-content in the home, blood–lead levels of members of the household, and long-run outcomes that have been linked to lead poisoning, just as we plan to do with the lead paint intervention (see Section 3.2).

As for data on our vendor, BRAC, all we need is simple accounting data: costs, quantities, revenues, etc. It would be nice to have the same data on the other sellers, these are notoriously difficult to measure properly. Instead, we may need to deploy secret buyers in order to observe marketing behavior and also collect samples to be tested for lead.

We will need to think carefully about how to design the surveys to not induce experimenter demand effects: it is easy to imagine how being asked about turmeric consumption and then having your blood tested for lead may cause you to put the pieces together, perhaps causing you to change your behavior.

[8]Also of some interest is the possibility of direct effects of the intervention on sellers: do sellers change behavior even if sellers do not?

4.4 Estimation: Discrete choice model

To capture consumer behavior, we will use the canonical discrete choice model described in Berry and Haile (2021), which will allow us to estimate heterogeneity in consumers' tastes for safety.

4.5 Alternative approaches

There are a few other ways we might consider trying to study the market for adulterated turmeric and shift it to a more salubrious equilibrium.

4.5.1 Randomize price

It may be interesting to cross-randomize along a third dimension: price, which may enable us to estimate a demand curve. However, there is reason to suspect we might conflate the price effect with another effect: signaling. While intuition implies that a new entrant will have a hard time making a dent on a market unless they price their product lower than the competitors — at least, as predicted in Tirole (1996), until they have built up a reputation — it is plausible that a higher price may actually *increase* demand in this setting. Without other means of signaling quality, in the No Testing arm the best option may be price: only a firm that is actually paying the cost of verifying that their turmeric is unadulterated would dare charging above the going market rate. On the other hand, adding this extra layer of randomization would give us an unwieldy number of combinations of treatment groups, we may be able to deploy the method developed in Banerjee et al. (2021) of wisely pooling treatment arms.

4.5.2 Intervene at the school level

There is one low-cost way to test for adulteration. Consumers who purchase whole turmeric roots instead of the ground powder can identify adulterated turmetic roots by snapping them in half: if the color of the inside does not match the color of the outside, that is highly indicative of adulteration during the polishing process (Forsyth et al., 2019a). Usually, all but the poorest households purchase ground turmeric, since the convenience far outweighs additional cost. However, there are some families who, fearing adulteration, pool their resources to purchase whole turmeric roots and then bring the roots to the miller themselves. Interestingly, they do this not to prevent adulteration with yellow dye, but adulteration with rice powder, which is harmless but dilutes the flavor and color of the turmeric; bringing the roots to the miller allows them to oversee the grinding process and prevent any funny business (Forsyth, 2022). There appears to be potential to leverage this demand for whole turmeric and promote awareness of adulteration with lead chromate.

A simple intervention would be to design an interactive lesson on turmeric adulteration to be delivered at schools: an expert would come to the school and explain the dangers of adulterate turmeric and the importance of avoiding it at all costs. Key to the effectiveness of the lesson would be an the interactive component: bringing samples, showing the difference between clean and adulterated roots, demonstrating how the children should test roots at markets, and allowing them to try it for themselves will make for a fun and memorable lesson. If successfully, households whose children attend treated schools will adopt the techniques and demand unadulterated turmeric, pushing the market equilibrium in the right direction.

5 Conclusion

In my quest to make the world a better place, I have yet to find a more promising way to do so than by studying lead poisoning in LMICs. Little is known about which sources are causing lead poisoning in LMICs, and few methods of eliminating those sources have been proposed, let alone tested. Considering the potential gains in quality of life to those who LMICs, this topic deserves much more attention than it is currently receiving. Perhaps some of that is due to the fact that the magnitude of those gains is extremely uncertain. Research that narrows those confidence bands on those gains, and that illuminates the path to achieving them, is sorely needed.

A Literature review

A.1 Physiology of Lead Poisoning

Lead primarily enters the body through dust particles, but may also be consumed directly through contaminated food or drink. Once in the body, lead is absorbed into the blood, where it has a half-life of about one month (Lidsky and Schneider, 2003). However, due to its chemical similarity to calcium, lead within the bloodstream can is easily absorbed by the bones. Lead deposits in bones will remain for decades, only to be released back into the bloodstream periodically (Barbosa et al., 2005; Hu et al., 2007). This allows ample opportunity for lead to jump the blood–brain barrier.

The harmful effects of lead in the brain are a result of the fact that calcium, in addition serving as the building block for our bones and teeth, is an important neurotransmitter (Lidsky and Schneider, 2003). Once lead is in the brain, due its resemblance to calcium, it effectively jumps the calcium queue. From there, lead directly interferes with communication between calcium-sensitive neurons, adding noise to the calcium signal (Minnema et al., 1988; Bressler and Goldstein, 1991).

Even more sinister is lead's effect on preventing proper brain development (Silbergeld, 1992). Lead prevents proper neuron growth and pruning, permanently altering the brain's architecture and typical dopamine and glutamine system development (Kern and Audesirk, 1995; Zawia and Harry, 1996; Patrick and Anderson, 2000; Zheng et al., 2001). In other words, lead causes stunting in brain growth. These effects seem to be concentrated especially in the pre-frontal cortex, the area of the brain responsible for executive control (Cecil et al., 2008; Funahashi and Andreau, 2013).

Clearly preventing proper brain development is a bigger problem for children than for adults, so preventing exposure in children is paramount. This is exacerbated by the fact that, conditional on lead exposure, children are more likely than adults to ingest lead — a result of their playing outside, crawling on the ground, putting their hands in their mouths, and consuming more food and drinks per body weight than adults (Leggett, 1993). However, even conditional on lead ingestion, it is a biological fact that a higher share of that ingested lead will be absorbed by the bloodstream in children than adults — especially for children with poor nutrition (Patrick, 2006).

A.2 Evidence from economics

For many years, the economics literature on lead was limited to (controlled) correlational studies at the population level. For example, Nevin (2000, 2007) and Mielke and Zahran (2012) track crime rates over time at the country or city level, and find that it rises in the 1970s, peaks in early 1990s, and declines thereafter, following the same pattern as the ambient lead levels but lagged by about 20 years. The lag is a result of a confluence two factors. The first factor is what's known as the age–crime curve, which describes the fact that the committing of crimes peaks around age 15–24 years, depending on the type of crime (Blumstein, 1995). The second factor, as discussed in section A.1, is that lead's effects are most severe in children. The result is that crime rates should rise as a function of the share of the population in peak crime-committing age who were exposed to lead in childhood, and fall as that share falls.

Jessica Wolpaw Reyes was the first to bring a quasi-experimental design to the study of the effect of lead poisoning within economics. Reyes (2007) used the market share of different gasoline octanes (i.e., regular versus premium) at the state level as an as-good-as-randomly assigned instrument for predicted lead exposure (since different octanes have different lead concentrations) and found that the phase-out of leaded gasoline in the US from 1975–85 in response to the Clean Air act caused violent crime to fall by 56%.

Reyes (2015a) uses a similar empirical strategy — with a large data set of nationally representative samples to be able to control for extremely covariates — and finds strong connections between between blood–lead levels at age 0–3 and various long-run outcomes, including an elasticity of around 0.1–1 for "behavior problems" (e.g., bullying, strong temper, hyperactivity, cheating, being easily confused, etc.) at age 4–12, "risky behavior" (e.g., teenage pregnancy, drug use, and alcohol use) and violence (e.g., attacking someone, getting arrested, etc.) at age 13–17. One more study that exploits variation in response to the phase-out of leaded gasolin in the US is Clay et al. (2021). Here, authors use proximity to state highways and an IV design that exploits noisy air pollution monitor measurements and find that the decrease in airborne lead over this period increased fertility by 0.14 children (6%).

Unlike gasoline used in cars, gasoline used in airplane is still permitted to contain lead in the US. The Environmental Protection Agency estimates that emission from airplanes are responsible for at least one half the the lead exposure in the United States. Zahran et al. (2017) use panel data to exploit variation in lead exposure as a result of proximity to airports in Michigan find social costs of about $10 per gallon, well in access of the $6 sticker price.

Reyes (2015b) employs a differences-in-differences design — with individual blood–

lead from Massachusetts, where all children under six are screened for lead, and school-by-year–level test scores from third and fourth grade in mathematics and English and language arts — and finds that every 1 percentage point reduction in share of children with blood–levels above 10 µg/dL causes a 0.2 percentage point reduction in share of children with unsatisfactory test score.

Aizer et al. (2018) takes advantage of the availability of data in Rhode Island that is a step up in quality from that of Massachusetts and links individual-level blood–lead levels to individual third grad test scores. They exploit a 1997 policy change that required landlords to ensure that their properties were lead-free, which disproportionately affected poor, African American neighborhoods: African American children's blood–lead levels fell . Instrumenting for blood–lead level using predicted probability of "lead-safe certificates" for rental properties, they find that a 1 µg/dL reduction in blood–lead level decreased the probability of being substantially below proficient in reading by 26%, and in math by (an insignificant) 13%. The greater impact on reading than on math was echoed by Sorensen et al. (2019), which uses variation causes by applications to the Center for Disease Control and Prevention for lead hazard reduction grants and find that a 1% decrease lead exposure in early childhood causes a 0.04 standard deviation increase in math scores and a 0.08 increase in reading.

Aizer and Currie (2019) return to their individual-level Rhode Island data but this time link it to individual-level data on school disciplinary records and juvenile arrests. They construct two different instrumental variables. The first exploits the fact that they have multiple blood–lead measurements for most children, and also measurements of the siblings within a household: instrumenting for noisy measurements eliminates attenuation bias in instrumental variables designs if the noise is purely random (Angrist and Krueger, 1999). Their second instrumental variable exploits the fact that ban on leaded gasoline eliminated the extremely tight correlation between proximity to a road and blood–lead levels. They find that a a 1 µg/dL reduction in blood–lead level decreases probability of school detentions by 27–74% in boys and suspensions by 6–9% in boys and girls.

Billings and Schnepel (2018) also exploit the fact that blood–lead tests are somewhat noisy, but use it in a design similar in spirit to a regression-discontinuity design. In North Carolina, two positive (i.e., $\geq$ 10 µg/dL) tests for a child triggers a multi-pronged intervention by the state including, among other things, interviews to identify lead sources in the home and referral to remediation programs to eliminate such sources. Children who test positive twice are compared those those who test positive only once, who should be (and are) similar on observable characteristics — except those in which we're interested here. The authors match blood–lead data to data on arrests and public school record for all children born 1990–1997 in one North Carolina county and find that, despite their much higher blood–lead levels at baseline, the intervention-eligible group had 48% fewer school suspensions, 33% fewer days of school missed, and 56% fewer arrests. There were also positive (though insignificant) point estimates on reading and math scores.

Hollingsworth and Rudik (2021) examine one exception to the ban on leaded gasoline that was grandfathered in for decades: automotive racing. The National Association of Stock Car Racing (NASCAR) and the Automotive Racing Club of America (ARCA) was permitted to and continued to use leaded gasoline until 2007 due to its effect on the efficiency of engines. They use a differences-in-differences design, comparing counties near to and far from racetracks before and after 2007 and find that races cause a 20 % jump in ambient lead measured by monitors within 50 miles within one week of the race (and that the effects dissipate within about one month), a 17% increase in percentage of children with elevated blood–lead levels aggregated at the county-year level, and 91 all-cause elderly deaths per 100,000 driven primarily by cardiovascular disease and ischemic heart disease.

Feigenbaum and Muller (2016) examine lead's effect on homicide in the US in the early 20th century. Before the widespread use of leaded gasoline, the largest source of lead exposure was water pipes: as of 1897, 54% of US cities used lead water pipes as opposed to iron pipes due to lead's durability and malleability, despite its higher cost. Ironically, these desirable characteristics meant that it was generally the wealthier, better-educated, more "health conscious" cities that were more likely to use lead pipes (Troesken, 2008). Using distance from a lead refinery as an instrument for a city's use of lead pipes, they find that lead pipes increase homicide rate by about 24%.

In an alternative specification, Feigenbaum and Muller (2016) find that when they include an interaction term between a lead pipe indicator and local water acidity, that explains a large part of the variation in homicides — a finding that should come as no surprise to those who remember the Flint, Michigan water crisis, which made common knowledge the fact that lead pipes will leak more lead particles into more acidic water — while the effect of acidity in cities that do no use lead pipes is actually goes in the opposite direction, meaning that more acidic water *per se* may actually be a benefit. This echos findings by Clay et al. (2014) Troesken (2008), which find a large, significant effect on infant mortality of acidic water in cities that use lead pipes.

B Collective reputation model

The situation can described with a collective reputation model as in Tirole (1996), which I will use here almost unaltered. Sellers can choose one of two actions:

1. Provide high-quality, unadulterated turmeric. This confers utility H to the buyer at cost C to the seller.

2. Provide low-quality, adulterated turmeric. This confers utility L at zero cost to the seller.

We assume that $H > L$ because, presumably, it is better to be not poisoned than to be poisoned. We make the (dubious) assumption that poisoned turmeric is better than no turmeric, so $L > 0$. We also assume $C > 0$ because it costs the seller some time and money to *not* adulterate because dye is actually cheaper per unit weight than turmeric. Finally, we assume that the high equilibrium — where all turmeric sold is unadulterated — brings more surplus than the low equilibrium: $H - C > L > 0$.

In the high equilibrium, buyers are willing to pay $p_H \leq H$, and sellers are willing to sell at $p_H \geq C$. Assume imperfect competition, so at equilibrium $H < p_H < C$. Then buyers get utility $H - p_H > H - C = 0$, and sellers make profit $\pi = p_H - C > C - C = 0$

However, a dishonest seller would rather advertise that they're selling high quality turmeric to be able to charge $H > p_H > C$, but then cheat and provide low quality turmeric and take home profit p_H rather than $p_H - C$. Buyers cannot *ex ante* differentiate between high- and low-quality turmeric, but they are aware of the sellers' temptation to cheat, so they instead offer $L \geq p_L \geq 0$, because $L - p_L > L - p_H$. Then buyers get utility $L - p_L$ and sellers make profit $\pi = p_L$. Even though $H - C > L$, we are stuck at the low equilibrium.

What are some ways to get out of the low equilibrium?

B.1 Perfect information

One for honest sellers to build reputation. An honest sellers should be willing to sell high quality turmeric at a loss if that endows it with reputation for selling high quality turmeric, which attracts future buyers and future profit. Therefore, so long as the number of buyers n is high enough that present discounted value of future profits $\Pi > L$, honest sellers will take home a profit in the long run. Assuming a Poisson death process λ (i.e., each period, any given seller has a probability λ of dying and being replaced by an identical seller) and a discount rate δ, the honest seller will provide high quality if

$$\Pi = \sum_{n=1}^{\infty}(1-\lambda\delta)^n(p_H - C) \geq L. \tag{11}$$

Sellers who cheat will be caught, have their reputation ruined, and never be able to command a price $p > L$. This implies scope for intervention. If the buyers have the ability to perfectly verify quality, then individual sellers can build reputation, and the market should shift to the high equilibrium. In the case of turmeric, this would be analogous with giving every buyer access to an XRF gun, which is unrealistic. We will have to find another solution.

B.2 Imperfect information

Let's assume that a fraction α of sellers is honest and never cheats if offered p_H, some fraction of sellers β is dishonest and always cheats no matter whether offered p_H or p_L. Let $\alpha + \beta = 1$.

Suppose that quality cannot always be detected. Instead, buyers have some probability of catching a cheating seller x_k, where k indexes the number of times this seller has cheated. The seller's reputation is binary: they are either "clean" (i.e., never caught up until this point) or "tainted" (i.e., caught at least once up until this point). This would represent a world where, for example, each transaction has some chance of being spot checked by a government agent with an XRF gun: if a seller is caught selling adulterate turmeric, their reputation is forever tainted. A buyer buying from a seller who has cheated k times in the past has a probability x_k of being aware of their tainted reputation. Assume that $x_0 = 0 < x_1 < x_2 < \cdots < 1$ (i.e., that a buyer is more likely to hear of the tainted reputation of a bigger cheater) and that $x_{k+1} - x_k < x_k - x_{k-1}$ (i.e., cheating has decreasing marginal reputation tainting). Putting these two assumptions together means that once a seller cheats once, they may as well keep cheating.

If their buyer knew what fraction of sellers where honest (α) and what fraction were dishonest (β) but had no information about this particular seller (i.e., $Y = 0$), their decision would be a simple one. The buyer would offer p_H if their expected payoff from offering p_H and hoping for high quality but running the risk of overpaying for low quality is greater than taking the sure thing of offering p_L.

$$\alpha(H - p_H) + \beta(L - p_H) > L - p_L \tag{12}$$

Allowing for some probability of detection, there are now two possible scenarios. If the buyer encounters a seller with a tainted reputation, they know the seller is dishonest and will adulterate no matter what, so the buyer offers p_L. If buyer encounters a seller with a clean reputation, there's still a chance that this is a dishonest seller who has yet to be caught. What are the chances of this? Well, a "newborn" seller cannot have cheated, so they have reputation $x_0 = 0$. According to our Poisson death process, they represent a fraction $1 - \lambda$ of the population. A seller who was "born" one period ago and cheated in their only opportunity to has reputation x_1 and makes up a fraction $(1 - \lambda)\lambda$ of the population. A seller born two periods ago and cheated both times has reputation x_2 and makes up a fraction $(1 - \lambda)\lambda^2$ of the population. Therefore, the average dishonest seller has the following probability of having an untainted reputation:

$$Y = (1 - \lambda) \sum_{k=0}^{\infty} \lambda^k (1 - x_k). \tag{13}$$

Whereas the fraction of sellers who have a tainted reputation is

$$Y - 1 = \sum_{k=0}^{\infty} (\lambda)^k x_{k+1}. \tag{14}$$

So now the population of dishonest can be divided up into three groups: honest sellers α, clean dishonest sellers βY, and tainted dishonest sellers $\beta(1 - Y)$. The fraction of untainted sellers that is honest is therefore $\frac{\alpha}{\alpha+\beta Y}$, and the fraction that is dishonest is $\frac{\beta Y}{\alpha+\beta Y}$. Therefore, conditional on meeting a seller with a clean reputation, the buyer's decision looks almost the same as it did before, except their payoffs are now weighted by these new probabilities.

$$\frac{\alpha}{\alpha + \beta Y}(H - p_H) + \frac{\beta Y}{\alpha + \beta Y}(L - p_H) > L - p_L. \tag{15}$$

In such a world, we have only two avenues for changing buyer choices. The first is to intervene on the demand side: for example, an information campaign that increases buyer awareness of adulteration could be modeled a number of ways (e.g., increasing

H or p_H or decreasing L or p_L), but regardless it could tilt the scales in favor of the high equilibrium. The other avenue is to increase probability of detection x_k: for example, if the government increased the frequency of its spot checks, that would decrease Y, which would also promote the high equilibrium.

B.3 Opportunistic sellers

This is, of course, an uninteresting world: we cannot do anything about the sellers, since their decision as to whether to offer adulterated turmeric is fixed. Now let us introduce a third type of seller, one that is "opportunistic" and will cheat when offered p_H only if they will profit in the long run. Now let $\alpha + \beta + \gamma = 1$. Let us model the opportunistic seller's decision. The opportunistic seller can choose to never cheat, in which case their payoff would be the same as that of an honest seller. If α is high enough to sustain the high equilibrium, then:

$$\Pi_H = \sum_{n=1}^{\infty} (\lambda\delta)^n (p_H - C) = \frac{p_H - C}{1 - \lambda\delta}. \tag{16}$$

If instead the seller cheats today and forever more, then their expected payoff is

$$\Pi_L = p_H + \lambda\delta p_H \left(\frac{1}{1 - \lambda\delta} - Z \right) + \lambda\delta(p_L + C)Z. \tag{17}$$

where Z is the probability of being caught, discounted back to the present:

$$Z = x_1 + \lambda\delta x_2 + (\lambda\delta)^2 x_3 + ... = \sum_{k=0}^{\infty} x_{k+1}(\lambda\delta)^k. \tag{18}$$

Setting $\Pi_H \geq \Pi_L$ and rearranging, we find that the opportunistic seller will not cheat if

$$\frac{C}{1 - \delta\lambda} \leq \delta\lambda(p_H - C - p_L)Z. \tag{19}$$

This implies the same interventions as before, such as an information campaign to decrease L or a improved government oversight to increase Z. In addition, we find an effect on the seller's decision from decreasing the cost of making the high-quality product: for example, if it becomes easier for the seller to verify that their suppliers are not adulterating the turmeric roots, then they'd be more likely to sell unadulterated ground turmeric.

Suppose this condition is satisfied, so the opportunistic sellers choose never to cheat. In that case, the buyer decision looks identical to how it did before, except now the opportunists γ get lumped in with the honest sellers α. The buyer will offer p_H if

$$\frac{\alpha+\gamma}{\alpha+\gamma+\beta Y}(H-p_H)+\frac{\beta Y}{\alpha+\gamma+\beta Y}(L-p_H)>L-p_L. \tag{20}$$

B.4 A shock

Bangladeshi turmeric farmers describe one particular season in which there was unusual flooding, which prevented the proper drying of the turmeric and robbed it of its attractive yellow color. They cite their inability to compete that season with turmeric from India, whose turmeric did not suffer the same fate, as the impetus of turmeric adulteration (Forsyth et al., 2019b). Let this unsuccessful season be modeled as a temporary shock to H and therefore also to p_H: due to its undesirable color, demand for unadulterated turmeric plummeted. Although the behavior of honest and dishonest sellers was unaffected, opportunists all cheated, sharply and permanently increasing their probability of being caught and having their reputation tainted forever. Although none of the parameters underlying equation 20 have changed, the decision-making rule has: opportunists are now cheaters, so should be lumped in with the dishonest sellers β, instead of the honest sellers α. If the share of opportunists γ is very high, this could flip the inequality:

$$\frac{\alpha}{\alpha+(\beta+\gamma)Y}(H-p_H)+\frac{(\beta+\gamma)Y}{\alpha+(\beta+\gamma)Y}(L-p_H)<L-p_L. \tag{21}$$

It is worth noting that equations 20 and 21 are not mutually exclusive. I will omit the details for the sake of brevity (see Tirole (1996)), but suffice it to say that given all of the same parameters, a temporary shock that alters the behavior of the opportunists can move the world from one equilibrium to the other.

B.5 Predictions of the model

Nearby Testing: Allowing buyers to bring their turmeric to a nearby government facility to be tested for lead should simply increase the probability that an adulterating firm gets caught doing so (Z). Allowing for some heterogeneity in the model, this should convince a share of sellers to stop adulterating, as their expected payoff from adulterating decreases as a result of the increased likelihood of having their reputation tainted.

Market Testing: Allowing the buyers to test their turmeric for lead at the same market at which they buy it in the first place will further increase the probability that an adulterating firm gets caught. In addition, the salience of such testing may serve as a shock to consumer preferences — *"If they are testing for lead, it must be really bad!"* — causing them to value adulterated turmeric less than they would otherwise. This should reduce their utility from low quality turmeric (L) and cause more sellers to stop adulterating.

Brand: Simply introducing BRAC-branded turmeric to the market increases the share of honest sellers (α). However, absent any means for the buyers to discover that that share has increased, this should have no direct effect on sellers' decision to adulterate. BRAC may be a big enough player to shift the supply curve outward significantly, but that will lower buyers' willingness to pay for high quality (p_H) and low quality (p_L) equally, so they effects cancel out.

Brand Claim: In addition to increasing the share of honest firms (α), introducing BRAC-branded turmeric that makes a salient claim that their turmeric is lead-free might increase valuation of high-quality turmeric (H) — *Why would BRAC talk about lead if it wasn't important?"* — which could increase willingness to pay (p_H), and cause some sellers to stop adulterating. However, this is conditional on consumers believing the claim: if they think a seller is lying, then the claim has no signaling value, and buyers will not change their behavior, so neither will sellers. In fact, this could even backfire: if consumers believe the claim even without proof, making them more willing to offer p_H, then some sellers will start making the claim while continuing to adulterate, increasing their market share.

Warranty: In addition to increasing the share of honest firms α and perhaps increasing utility from high quality H, introducing BRAC-branded turmeric that makes a salient claim that their turmeric is lead-free and also that their turmeric comes with a more-than-your-money-back guarantee may also increase the frequency with which buyers test their turmeric for lead, increasing the probability that a seller that adulterates will be caught (Z). Furthermore, it increases the cost to the seller of adulterating (because if they do so, they risk having to refund buyers), which effectively decreases the cost of *not* adulterating (C), which tilts the balance in favor of even more sellers choosing not to adulterate.

C Other avenues for research

Leaded gasoline, lead paint, and adulterated turmeric are far from the only sources of lead poisoning in LMICs. It will certainly be worthwhile to study other likely sources in order to better understand how reduce the prevalence of lead poisoning.

C.1 Informal Recycling of Used Lead-Acid Batteries

Approximately 85% of the world's lead production goes in to car batteries (Tür et al., 2016), but, at least in HICs, this lead is processed in highly regulated facilities that cause minimal exposure. This includes when the batteries are recycled: because lead is infinitely recyclable, the majority of the inputs to lead-acid batteries can be recovered from used lead-acid batteries (ULABs), minimizing the need to mine for new lead (Rees and Fuller, 2020). Unfortunately, the regulations around ULAB recycling are not nearly as strict, nor as effective in LMICs. For example, when a policy change shifted a huge share of ULAB recycling from the US to Mexico, this caused serious adverse health effects in the neighborhoods immediately surrounding the Mexican ULAB recycling facilities (Tanaka et al., 2022).

Unfortunately, most ULAB recycling in LMICs is much, much more dangerous than the ones in Mexico. Typically, batteries are rounded up from various sources (mechanics, scrap yards, etc.), chopped open with machetes, and then separated into its various components: the plastic gets tossed into a trash heap, the liquid acid is dumped indiscriminately on the ground, and the lead is thrown into a furnace to be melted into ingots (Gottesfeld and Pokhrel, 2011; Belay et al., 2015). The potential for lead poisoning is obvious: the plastic could exposure children playing near the trash heaps, the liquid could contaminate drinking water, and the furnace could blanking the entire surrounding area in lead soot. One study estimates that informal ULAB recycling is responsible for between 125,000–1,613,000 DALYs in 2013 alone, making it the largest single contributor to lead poisoning in the world (Ericson et al., 2016).

One possible intervention would be to take inspiration from the Öko Institute (Manhart et al., 2014), which transported an entire ship's worth of ULABs from Ghana to Germany to be properly recycled. An intervention in which we create a market to which people who currently work as informal ULAB recyclers sell whole (i.e., yet-to-be-recycled) ULABs to reputable battery manufacturers in HIC, where the recycling can be done safely, seems like a promising way to prevent lead poisoning in places where this is a problem — primarily Sub-Saharan Africa (Gottesfeld et al., 2018).

C.2 Ceramics

Lead glazes are used in traditional pottery in Mexico and elsewhere in Latin America (Charlton, 1976). When done properly, the ceramics should be safe for use in cooking. However, it has been discovered that the kilns used by many traditional potters does not get hot enough to completely "vitrify" (i.e., solidify) the glaze, allowing for lead to leach from the glaze into the food when cooking (Caravanos et al., 2014). An intervention in which potters are randomized at the market or community level into a program to provides more powerful kilns and training on how to use them seems promising.

C.3 Ship-breaking

Ship-breaking is the recycling of the large ships (e.g., oil tankers) into scrap metal. In HICs, this can be done relatively safely, but it is prohibitively expensive. Instead, ships are typically recycled unskilled laborers in LMICs, where ship-breaking is an extremely dangerous profession, not just because of the high probability of injury and death due to falls and other accidents, but also because of the unfathomable amount of pollution this process creates (Sahu, 2014; Demaria, 2010; Nøst et al., 2015). The vast majority of it (around 90% of annual tonnage recycled) occurs at three ports: Alang in India, Chittagong in Bangladesh, and Gadani in Pakistan (Fairtrade, 2020). This was not always the case: in 2019, China suddenly stopped accepting ships for recycling if they did not fly Chinese flags. Up until that point, China had be handling about 20% of annual recycled tonnage, so this policy change was a sizable shock to the market, which could serve as a useful source of exogenous variation. If I can a data set with high enough spatial and temporal resolution — perhaps hospital birth records, *à la* Tanaka et al. (2022) — I can tease apart a causal effect of shipbreaking in these areas. Estimating the first stage is fairly straightforward: use satellite data to simply count ships.

References

Abbott, E. (1877). *Long Look House: A Book for Boys and Girls*. Noyes, Snow and Company, Boston.

Aizer, A. and Currie, J. (2019). Lead and Juvenile Delinquency: New Evidence from Linked Birth, School, and Juvenile Detention Records. *The Review of Economics and Statistics*, 101(4):575–587.

Aizer, A., Currie, J., Simon, P., and Vivier, P. (2018). Do Low Levels of Blood Lead Reduce Children's Future Test Scores? *American Economic Journal: Applied Economics*, 10(1):307–341.

Allcott, H., Lockwood, B. B., and Taubinsky, D. (2019). Should We Tax Sugar-Sweetened Beverages? An Overview of Theory and Evidence. *Journal of Economic Perspectives*, 33(3):202–227.

Angrist, J. D. (1990). Lifetime Earnings and the Vietnam Era Draft Lottery: Evidence from Social Security Administrative Records. *The American Economic Review*, 80(3):313–336. Publisher: American Economic Association.

Angrist, J. D. and Krueger, A. B. (1999). Empirical Strategies in Labor Economics. In *Handbook of Labor Economics*, volume 3, pages 1277–1366. Elsevier.

Attina, T. M. and Trasande, L. (2013). Economic Costs of Childhood Lead Exposure in Low- and Middle-Income Countries. *Environmental Health Perspectives*, 121(9):1097–1102. Publisher: Environmental Health Perspectives.

Bai, J. (2016). Melons as Lemons: Asymmetric Information, Consumer Learning and Seller Reputation. Natural Field Experiments 00540, The Field Experiments Website.

Bai, J., Gazze, L., and Wang, Y. (2022). Collective Reputation in Trade: Evidence from the Chinese Dairy Industry. *The Review of Economics and Statistics*, 104(6):1121–1137.

Banerjee, A., Chandrasekhar, A. G., Dalpath, S., Duflo, E., Floretta, J., Jackson, M. O., Kannan, H., Loza, F. N., Sankar, A., Schrimpf, A., and Shrestha, M. (2021). Selecting the Most Effective Nudge: Evidence from a Large-Scale Experiment on Immunization.

Barbosa, F., Tanus, S. J. E., Gerlach, R. F., and Parsons, P. J. (2005). A Critical Review of Biomarkers Used for Monitoring Human Exposure to Lead: Advantages, Limitations, and Future Needs. *Environmental Health Perspectives*, 113(12):1669–1674. Publisher: Environmental Health Perspectives.

Bartos, V., Bauer, M., Chytilová, J., and Levely, I. (2021). Psychological Effects of Poverty on Time Preferences. *The Economic Journal*, 131(638):2357–2382.

Baumer, E. P., Rosenfeld, R., and Wolff, K. T. (2012). Expanding the scope of research on recent crime trends. *Annotation*.

Belay, M., Belay, A., and Genet, Z. (2015). Safety practices and awareness of lead acid battery recyclers in Addis Ababa, Ethiopia. *Addis Ababa: Pesticide Action Nexus Association*.

Berry, S. T. and Haile, P. A. (2021). Chapter 1 - Foundations of demand estimation. In Ho, K., Hortaçsu, A., and Lizzeri, A., editors, *Handbook of Industrial Organization*, volume 4 of *Handbook of Industrial Organization, Volume 4*, pages 1–62. Elsevier.

Billings, S. B. and Schnepel, K. T. (2018). Life after Lead: Effects of Early Interventions for Children Exposed to Lead. *American Economic Journal: Applied Economics*, 10(3):315–344.

Björkman Nyqvist, M., Svensson, J., and Yanagizawa-Drott, D. (2022). Can Good Products Drive Out Bad? A Randomized Intervention in the Antimalarial Medicine Market in Uganda. *Journal of the European Economic Association*, 20(3):957–1000.

Blumstein, A. (1995). Youth Violence, Guns, and the Illicit-Drug Industry. *The Journal of Criminal Law and Criminology (1973-)*, 86(1):10.

Bressler, J. P. and Goldstein, G. W. (1991). Mechanisms of lead neurotoxicity. *Biochemical Pharmacology*, 41(4):479–484.

Caravanos, J., Dowling, R., María Téllez-Rojo Dra, M., Cantoral, A., Kobrosly, R., Estrada, D., Orjuela, M., Gualtero, S., Ericson MSc, B., Rivera, A., and Fuller, R. (2014). Blood Lead Levels in Mexico and Pediatric Burden of Disease Implications. *Annals of Global Health*, 80(4):269.

Cecil, K. M., Brubaker, C. J., Adler, C. M., Dietrich, K. N., Altaye, M., Egelhoff, J. C., Wessel, S., Elangovan, I., Hornung, R., Jarvis, K., and Lanphear, B. P. (2008). Decreased Brain Volume in Adults with Childhood Lead Exposure. *PLOS Medicine*, 5(5):e112. Publisher: Public Library of Science.

Charlton, T. H. (1976). Contemporary Central Mexican Ceramics: A View from the Past. *Man*, 11(4):517.

Clay, K., Portnykh, M., and Severnini, E. (2021). Toxic Truth: Lead and Fertility. *Journal of the Association of Environmental and Resource Economists*, 8(5):975–1012. Publisher: The University of Chicago Press.

Clay, K., Troesken, W., and Haines, M. (2014). Lead and Mortality. *Review of Economics and Statistics*, 96(3):458–470.

Cohen, J. and Dupas, P. (2010). Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment*. *The Quarterly Journal of Economics*, 125(1):1–45.

Coulter, L. (2022). Conversation with an industry expert on lead paint in low- and middle-income countries.

Cowell, W., Ireland, T., Vorhees, D., and Heiger-Bernays, W. (2017). Ground Turmeric as a Source of Lead Exposure in the United States. *Public Health Reports*, 132(3):289–293. Publisher: SAGE Publications Inc.

Demaria, F. (2010). Shipbreaking at Alang–Sosiya (India): An ecological distribution conflict. *Ecological Economics*, 70(2):250–260.

Dignam, T., Kaufmann, R. B., LeStourgeon, L., and Brown, M. J. (2019). Control of Lead Sources in the United States, 1970-2017: Public Health Progress and Current Challenges to Eliminating Lead Exposure. *Journal of public health management and practice : JPHMP*, 25(Suppl 1 LEAD POISONING PREVENTION):S13–S22.

Donohue, J. and Levitt, S. (2019). The Impact of Legalized Abortion on Crime over the Last Two Decades. Technical Report w25863, National Bureau of Economic Research, Cambridge, MA.

Douglas, R. (2021). Detecting Trucks in East Africa | Data Science Campus.

Duflo, E. (2001). Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment. *American Economic Review*, 91(4):795–813.

Ericson, B., Hu, H., Nash, E., Ferraro, G., Sinitsky, J., and Taylor, M. P. (2021). Blood lead levels in low-income and middle-income countries: a systematic review. *The Lancet Planetary Health*, 5(3):e145–e153.

Ericson, B., Landrigan, P., Taylor, M. P., Frostad, J., Caravanos, J., Keith, J., and Fuller, R. (2016). The Global Burden of Lead Toxicity Attributable to Informal Used Lead-Acid Battery Sites. *Annals of Global Health*, 82(5):686–699.

Eschnauer, H. R. and Stoeppler, M. (1992). Wine – an enological specimen bank. In *Techniques and Instrumentation in Analytical Chemistry*, volume 12, pages 49–71. Elsevier.

Fairtrade, M. (2020). Majority of global ship breaking takes place at Bangladesh, India, Pakistan.

Feigenbaum, J. J. and Muller, C. (2016). Lead exposure and violent crime in the early twentieth century. *Explorations in Economic History*, 62:51–86.

Forsyth, J. (2022). Conversation with an industry expert on turmeric adulteration in South Asia.

Forsyth, J. E., Nurunnahar, S., Islam, S. S., Baker, M., Yeasmin, D., Islam, M. S., Rahman, M., Fendorf, S., Ardoin, N. M., Winch, P. J., and Luby, S. P. (2019a). Turmeric means "yellow" in Bengali: Lead chromate pigments added to turmeric threaten public health across Bangladesh. *Environmental Research*, 179:108722.

Forsyth, J. E., Saiful Islam, M., Parvez, S. M., Raqib, R., Sajjadur Rahman, M., Marie Muehe, E., Fendorf, S., and Luby, S. P. (2018). Prevalence of elevated blood lead levels among pregnant women and sources of lead exposure in rural Bangladesh: A case control study. *Environmental Research*, 166:1–9.

Forsyth, J. E., Weaver, K. L., Maher, K., Islam, M. S., Raqib, R., Rahman, M., Fendorf, S., and Luby, S. P. (2019b). Sources of Blood Lead Exposure in Rural Bangladesh. *Environmental Science & Technology*, 53(19):11429–11436. Publisher: American Chemical Society.

Funahashi, S. and Andreau, J. M. (2013). Prefrontal cortex and neural mechanisms of executive function. *Journal of Physiology-Paris*, 107(6):471–482.

Gazze, L. (2022). Hassles and Environmental Health Screenings: Evidence from Lead Tests in Illinois. *Journal of Human Resources*, page 0221. Publisher: University of Wisconsin Press.

Gazze, L., Persico, C., and Spirovska, S. (2021). The Long-Run Spillover Effects of Pollution: How Exposure to Lead Affects Everyone in the Classroom.

Gottesfeld, P. and Pokhrel, A. (2011). Review: Lead Exposure in Battery Manufacturing and Recycling in Developing Countries and Among Children in Nearby Communities. *Journal of occupational and environmental hygiene*, 8:520–32.

Gottesfeld, P., Were, F. H., Adogame, L., Gharbi, S., San, D., Nota, M. M., and Kuepouo, G. (2018). Soil contamination from lead battery manufacturing and recycling in seven African countries. *Environmental Research*, 161:609–614.

Hamory, J., Miguel, E., Walker, M., Kremer, M., and Baird, S. (2020). Twenty Year Economic Impacts of Deworming. Technical Report w27611, National Bureau of Economic Research, Cambridge, MA.

Higney, A., Hanley, N., and Moro, M. (2022). The lead-crime hypothesis: A meta-analysis. *Regional Science and Urban Economics*, 97:103826.

Hodge, A. T. (1981). Vitruvius, Lead Pipes and Lead Poisoning. *American Journal of Archaeology*, 85(4):486–491.

Hollingsworth, A. and Rudik, I. (2021). The Effect of Leaded Gasoline on Elderly Mortality: Evidence from Regulatory Exemptions. *American Economic Journal: Economic Policy*, 13(3):345–373.

Hu, H., Shih, R., Rothenberg, S., and Schwartz, B. S. (2007). The Epidemiology of Lead Toxicity in Adults: Measuring Dose and Consideration of Other Methodologic Issues. *Environmental Health Perspectives*, 115(3):455–462. Publisher: Environmental Health Perspectives.

Jashemski, W. F. and Meyer, F. G. (2002). *The natural history of Pompeii*. Cambridge University Press, Cambridge ; New York.

Kern, M. and Audesirk, G. (1995). Inorganic lead may inhibit neurite development in cultured rat hippocampal neurons through hyperphosphorylation. *Toxicology and Applied Pharmacology*, 134(1):111–123.

Lacey, M. (2004). Belatedly, Africa Is Converting to Lead-Free Gasoline. *The New York Times*.

Lanphear, B. P., Hornung, R., Khoury, J., Yolton, K., Baghurst, P., Bellinger, D. C., Canfield, R. L., Dietrich, K. N., Bornschein, R., Greene, T., Rothenberg, S. J., Needleman, H. L., Schnaas, L., Wasserman, G., Graziano, J., and Roberts, R. (2005). Low-Level Environmental Lead Exposure and Children's Intellectual Function: An International Pooled Analysis. *Environmental Health Perspectives*, 113(7):894–899. Publisher: Environmental Health Perspectives.

Leggett, R. W. (1993). An age-specific kinetic model of lead metabolism in humans. *Environmental Health Perspectives*, 101(7):598–616. Publisher: Environmental Health Perspectives.

Lessler, M. A. (2006). Lead and Lead Poisoning from Antiquity to Modern Times. Publisher: Center for Biological Diversity.

Levitt, S. D. (2004). Understanding Why Crime Fell in the 1990s: Four Factors that Explain the Decline and Six that Do Not. *Journal of Economic Perspectives*, 18(1):163–190.

Lidsky, T. I. and Schneider, J. S. (2003). Lead neurotoxicity in children: basic mechanisms and clinical correlates. *Brain*, 126(1):5–19.

Manhart, A., Schleicher, T., and Degreif, S. (2014). Global Circular Economy of Strategic Metals–the Best-of-two-Worlds Approach (Bo2W). *Oeko-Institut eV*.

Mielke, H. W. and Zahran, S. (2012). The urban rise and fall of air lead (Pb) and the latent surge and retreat of societal violence. *Environment International*, 43:48–55.

Minnema, D. J., Michaelson, I. A., and Cooper, G. P. (1988). Calcium efflux and neurotransmitter release from rat hippocampal synaptosomes exposed to lead. *Toxicology and Applied Pharmacology*, 92(3):351–357.

Muralidharan, K. and Prakash, N. (2017). Cycling to School: Increasing Secondary School Enrollment for Girls in India. *American Economic Journal: Applied Economics*, 9(3):321–350.

Muralidharan, K., Romero, M., and Wüthrich, K. (2019). Factorial Designs, Model Selection, and (Incorrect) Inference in Randomized Experiments.

Nevin, R. (2000). How Lead Exposure Relates to Temporal Changes in IQ, Violent Crime, and Unwed Pregnancy. *Environmental Research*, 83(1):1–22.

Nevin, R. (2007). Understanding international crime trends: The legacy of preschool lead exposure. *Environmental Research*, 104(3):315–336.

Nøst, T. H., Halse, A. K., Randall, S., Borgen, A. R., Schlabach, M., Paul, A., Rahman, A., and Breivik, K. (2015). High Concentrations of Organic Contaminants in Air from Ship Breaking Activities in Chittagong, Bangladesh. *Environmental Science & Technology*, 49(19):11372–11380. Publisher: American Chemical Society.

Patrick, G. W. and Anderson, W. J. (2000). Dendritic alterations of cerebellar Purkinje neurons in postnatally lead-exposed kittens. *Developmental Neuroscience*, 22(4):320–328.

Patrick, L. (2006). Lead toxicity, a review of the literature. Part 1: Exposure, evaluation, and treatment. *Alternative Medicine Review: A Journal of Clinical Therapeutic*, 11(1):2–22.

Prasad, S. and Aggarwal, B. B. (2011). *Turmeric, the Golden Spice*. CRC Press/Taylor & Francis. Publication Title: Herbal Medicine: Biomolecular and Clinical Aspects. 2nd edition.

Rambachan, A. and Roth, J. (2022). A More Credible Approach to Parallel Trends. *Review of Economic Studies*.

Rees, N. and Fuller, R. (2020). *The toxic truth: children's exposure to lead pollution undermines a generation of future potential*. UNICEF.

Reyes, J. W. (2007). Environmental Policy as Social Policy? The Impact of Childhood Lead Exposure on Crime. *The B.E. Journal of Economic Analysis & Policy*, 7(1). Publisher: De Gruyter.

Reyes, J. W. (2015a). Lead Exposure and Behavior: Effects on Antisocial and Risky Behavior Among Children and Adolescents. *Economic Inquiry*, 53(3):1580–1605. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ecin.12202.

Reyes, J. W. (2015b). Lead Policy and Academic Performance: Insights from Massachusetts. *Harvard Educational Review*, 85(1):75–107.

Ritchie, H. (2022). How the world eliminated lead from gasoline.

Rosales-Rimache, J., Chavez-Ruiz, M., Inolopú-Cucche, J., Rabanal-Sanchez, J., Rueda-Torres, L., and Sanchez-Holguin, G. (2022). Leadcare® II Comparison with Graphite Furnace Atomic Absorption Spectrophotometry for Blood Lead Measurement in Peruvian Highlands. *Indian Journal of Clinical Biochemistry*.

Rosenfeld, R. and Fornango, R. (2007). The Impact of Economic Conditions on Robbery and Property Crime: The Role of Consumer Sentiment*. *Criminology*, 45(4):735–769. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1745-9125.2007.00096.x.

Sahu, G. (2014). Workers of Alang-Sosiya: A Survey of Working Conditions in a Ship-Breaking Yard, 1983-2013. *Economic and Political Weekly*, 49(50):52–59. Publisher: Economic and Political Weekly.

Silbergeld, E. K. (1992). Mechanisms of lead neurotoxicity, or looking beyond the lamppost. *The FASEB Journal*, 6(13):3201–3206. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1096/fasebj.6.13.1397842.

Sorensen, L. C., Fox, A. M., Jung, H., and Martin, E. G. (2019). Lead exposure and academic achievement: evidence from childhood lead poisoning prevention efforts. *Journal of Population Economics*, 32(1):179–218.

Stasik, M., Byczkowska, Z., Szendzikowski, S., and Fiedorczuk, Z. (1969). Acute tetraethyllead poisoning. *Archiv für Toxikologie*, 24(4):283–291.

Stiles, K. M. and Bellinger, D. C. (1993). Neuropsychological correlates of low-level lead exposure in school-age children: a prospective study. *Neurotoxicology and Teratology*, 15(1):27–35.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Tanaka, S., Teshima, K., and Verhoogen, E. (2022). North-South Displacement Effects of Environmental Regulation: The Case of Battery Recycling. *American Economic Review: Insights*, 4(3):271–288.

Tcherni-Buzzeo, M. (2019). The "Great American Crime Decline": Possible Explanations. In Krohn, M. D., Hendrix, N., Penly Hall, G., and Lizotte, A. J., editors, *Handbook on Crime and Deviance*, Handbooks of Sociology and Social Research, pages 309–335. Springer International Publishing, Cham.

Tirole, J. (1996). A Theory of Collective Reputations (with applications to the persistence of corruption and to firm quality). *The Review of Economic Studies*, 63(1):1–22.

Troesken, W. (2008). Lead Water Pipes and Infant Mortality at the Turn of the Twentieth Century. *Journal of Human Resources*, 43(3):553–575. Publisher: University of Wisconsin Press.

Tür, M., Manhart, A., and Schleicher, T. (2016). Generation of used lead-acid batteries in Africa–estimating the volumes. *Freiburg: Oeko-Institut eV. Available at: https://www. econet. international/fileadmin/user_upload/ULAB_Generation_African_Countries_final_20160411. pdf (accessed 6 December 2019).*

Ulmer, J. T. and Steffensmeier, D. (2014). The age and crime relationship: Social variation, social explanations. In *The Nurture Versus Biosocial Debate in Criminology*, pages 377–396. SAGE Publications Inc.

WHO (2022). Lead poisoning.

Wiebe, S. A., Sheffield, T. D., and Espy, K. A. (2012). Separating the Fish From the Sharks: A Longitudinal Study of Preschool Response Inhibition. *Child Development*, 83(4):1245–1261. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8624.2012.01765.x.

Zahran, S., Iverson, T., McElmurry, S. P., and Weiler, S. (2017). The Effect of Leaded Aviation Gasoline on Blood Lead in Children. *Journal of the Association of Environmental and Resource Economists*, 4(2):575–610. Publisher: The University of Chicago Press.

Zawia, N. H. and Harry, G. J. (1996). Developmental exposure to lead interferes with glial and neuronal differential gene expression in the rat cerebellum. *Toxicology and Applied Pharmacology*, 138(1):43–47.

Zheng, W., Lu, Y. M., Lu, G. Y., Zhao, Q., Cheung, O., and Blaner, W. S. (2001). Transthyretin, thyroxine, and retinol-binding protein in human cerebrospinal fluid: effect of lead exposure. *Toxicological Sciences: An Official Journal of the Society of Toxicology*, 61(1):107–114.