Emissions Demand Within and Across Sectors: Identifying the Effective Channels*

J. Scott Holladay,† Lawrence D. LaPlue, III‡

Abstract

We develop a decomposition to show how aggregate pollution emission intensity is driven by four establishment-level channels: entry, exit, reallocation of resources between survivors, and within-establishment adjustment of emissions intensity. Using a panel of establishment-level output and toxic pollution emissions data for U.S. manufacturers, we first empirically decompose changes in the three channels typically presented in the literature: scale (output), composition (sector market share), and aggregate technique (emissions intensity). We find that toxic pollution emissions from U.S. manufacturing have fallen by more than fifty percent driven largely by reductions in emissions intensity. We then decompose the aggregate emissions intensity channel into the four constituent channels. Our results show that reallocation towards cleaner firms and establishments as well as cleanup with surviving establishments are both important channels driving down aggregate emissions. Roughly three-quarters of the reduction in aggregate emissions intensity is driven by reduced emissions intensity within surviving establishments and the remaining quarter due to reallocation of production towards cleaner establishments. The cleanup occurs across all media (air and water) of emissions.

JEL Classification: C10, N50, Q56

Keywords: Emissions Decomposition, Establishment Entry and Exit, Aggregate Emissions, Emissions Intensity

1 Introduction

Over the past fifty years pollution emissions from U.S. manufacturers have been declining steadily, while real output has increased. Various factors driving this fall in emissions have

*We would like to thank Georg Schaur, Kip Viscusi, and Joel Rodrigue for helpful discussions.

†Department of Economics, University of Tennessee, Knoxville | jhollad3@utk.edu

‡Department of Economics, New Mexico State University, Las Cruces | llaplue@nmsu.edu

been debated in the economics literature, with environmental regulation, international trade and changing composition of output all considered as potential candidates. A general consensus has emerged that within sector reductions in pollution emissions intensity, termed technique effects, have played the largest role in the cleanup. The question of why emissions intensity has been falling remains difficult to answer. The cleanup could be due to adjustment of production techniques and investment in abatement on the part of existing establishments, or resources being reallocated, within sectors, to relatively less emissions intense establishments, or due to the exit of emissions-intense incumbents and the entry of relatively clean establishments.[1]

In this paper we apply a theoretical approach to decompose changes in within sector emissions intensity, the aggregate technique effect, into four within-sector channels. The first three: entry, exit, and relative changes in output by surviving establishments, can be combined into what we term the "reallocation effect." The fourth channel driving within-sector changes in emissions intensity that we identify is the change due to within-establishment reductions in emissions intensity, which is the "technique effect." We then use matched data on manufacturing establishment characteristics and toxic pollution emissions to identify the relative contributions of each of those four channels to the cleanup. The results indicate that the reallocation effect has been an important factor driving down emissions as output share within sectors has tended to shift towards less pollution intensive establishments. In addition, large declines in within-establishment pollution intensity, the technique effect, have also contributed to substantial declines in manufacturing emissions.

Copeland and Taylor (2003) develops a model that separates changes in pollution emissions into scale, composition and technique effects.[2] The scale effect is the increase in pollution associated with increased output. The composition effect is the change in pollution emissions due to shifts in economic activity towards or away from emissions-intense sectors. If cleaner sectors grow more quickly than dirtier sectors, the composition effect will drive down aggregate emissions per unit of output. Copeland and Taylor (2003) model the technique effect as the reduction in pollution that occurs as a polluting sector employs cleaner production techniques and lowers its aggregate emissions intensity in response to environmental regulation. In practice, any reduction in within sector emissions intensity is typically referred to as a technique effect, no matter the reason for the clean up.

Empirical estimates of these channels consistently find the technique effect plays the most important role in declining emissions. Levinson (2009) evaluates the impact of interna-

[1] Other possibilities, discussed by Cherniwchan et al. (2017), include changes in output product mix within an industry, or partial offshoring of intermediate aspects of production that are more emissions intense.

[2] Copeland and Taylor (2003) formally models the channels introduced in Grossman and Krueger (1993) focusing on the relationship between trade and the environment.

tional trade and decomposes changes in emissions of four air pollutants regulated under the Clean Air Act, known as criteria pollutants, into scale, composition, and aggregate technique channels. He finds that more than three-quarters of the observed reduction in emissions is attributable to the aggregate technique effect.[3] Levinson (2015) measures the technique effect directly, using dis-aggregated data, and finds that the technique effect accounts for as much as ninety percent of the total reduction in criteria air pollution emissions. Shapiro and Walker (2016) find that the majority of reduction in criteria air pollution emissions is driven by within product changes in pollution emissions and the observed clean up is consistent with a near doubling of the level of environmental regulation stringency between 1990 and 2008.[4]

We contribute to this literature in two ways. First, we apply a theoretical framework to decompose changes in the aggregate technique effect to cleanly identify the channels that drive observed reductions in sector emissions intensity. The decomposition allows us to isolate reductions in emissions intensity driven by existing establishments becoming cleaner, or output being reallocated from relatively dirty to relatively clean sub-industries, products, or firms. Second, we implement the decomposition using a new data set that includes matched establishment characteristics and toxic pollution emissions.[5] The existing literature focuses largely on emissions of criteria air pollutants, the six chemicals regulated under the Clean Air Act. We contribute to this literature by documenting a similar decline in pollution emissions for a broad set of toxic chemicals released across multiple release channels[6] and often subject to a different regulatory environment.

We begin by laying out a decomposition of aggregate manufacturing emissions into the three channels reported in the existing literature: scale, composition and technique. We extend this decomposition by modifying a method developed by Melitz and Polanec (2015)[7] to show how the aggregate technique effect is driven by four establishment-level channels: reallocation of resources among entering, exiting, and surviving establishments, and within-establishment adjustment to production techniques and emissions intensity. This permits us

[3]The paper acknowledges the possibility of a within-sector technique effect that could be covering up a composition effect that cannot be identified with 4-digit Standard Industrial Classification (SIC) data. We take advantage of more dis-aggregated data to examine this directly.

[4]They perform and updated version of the decomposition developed in Levinson (2009), by using more dis-aggregated 5-digit SIC data to look within "products," and not merely within 4-digit SIC "industries."

[5]Cherniwchan et al. (2017) develop a closely related decomposition that highlights the same channels (and others), potentially driving the aggregate technique effect. However, they do not develop a way to link the various channels together empirically.

[6]Toxic Release Inventory (TRI) reporters must document the media of release for their chemicals, e.g. to water, to air, to landfills or offsite incineration, or recycled.

[7]Melitz and Polanec (2015) focuses on decomposing industry productivity changes rather than emissions intensity.

to directly link establishment-level changes to aggregate emissions outcomes.

Identifying the importance of the establishment level channels of reduced emissions intensity requires information on output and pollution at the establishment level. We match pollution data from the EPA's Toxic Release Inventory (TRI) to establishment characteristics from the National Establishment Time Series (NETS) to create an unbalanced panel consisting of establishment-level observations of polluters between 1990-1999. Though several other papers in the literature have relied on Census of Manufacturing data for establishment characteristics and the National Emissions Inventory for air pollution emissions of chemicals regulated under the Clean Air Act, more recent papers have used a similar combination of NETS and TRI data to provide detailed analysis of the environmental effects of firm exporting (see Cui et al. (2016)) and the NAFTA trade liberalization (see Cherniwchan (2017)).[8] Our analysis provides additional evidence that the reduction in manufacturing pollution emissions extends beyond the criteria air pollutants to include toxic pollution emissions and emissions in other media besides air. These results suggest either spillovers from Clean Air Act regulation or that some other force underlies the cleanup of both types of pollution emissions.

Identifying the channel responsible for the cleanup is important. Understanding whether exit of polluting establishments or existing establishments changing their production process helps us identify the impacts of both trade and environmental regulation, a question of considerable debate in the literature.[9] If the cleanup is driven largely by the exit of polluting establishments it would lend credence to concerns that environmental regulation might harm the competitiveness of U.S. manufacturing. If the cleanup is driven by within sector reallocation to the cleanest establishments that might lead to market power concerns in those industries.[10] In addition, if other market forces are leading to the cleanup it complicates our evaluation of the effectiveness of environmental regulation.

Our results suggest that the main driver of the substantial clean up in U.S. manufacturing is reducing emissions intensity within surviving establishments. However, reallocation across firms and establishments is also important. While the emissions intensities of new entrants tend, on average, to be quite similar to those of existing firms, exiting firms tend to be dirtier than the firms they leave behind. In addition, intra-industry reallocation of output towards cleaner establishments has been an important source of cleanup during our sample period. When analyzing the channels across media of release, we document an overall decline in

[8] Cherniwchan (2017) estimates that the effects of NAFTA account, on average, for nearly two-thirds of the reductions in emissions documented by Levinson (2009) and Levinson (2015).

[9] See Jaffe et al. (1995) for a summary of the early literature and Greenstone et al. (2012).

[10] Ryan (2012) demonstrates how environmental regulation can lead to competitiveness inhibiting levels of regulation for example.

airborne emissions but a decline approximately half as large in aggregate emissions released to water and public treatment works, consistent with recent work by Gibson (2016) who documents a tendency of firms to substitute away from airborne emissions towards water when subjected to Clean Air Act regulations.[11]

The paper proceeds as follows. Section 2 introduces the theoretical decomposition of changes in pollution emissions first to the scale, composition and aggregate technique effects. We then demonstrate how the aggregate technique effect can be further decomposed into entry, exit, reallocation and within-establishment technique channels. Section 3 introduces the data and the empirical approach and presents the decomposition results. Section 4 explores emissions across media and section 5 presents several robustness checks. Section 6 concludes.

2 Decomposing Emissions Changes

This section briefly describes the standard decomposition performed in the literature and then lays out our additional decomposition of the aggregate technique effect into establishment-level channels.

2.1 Within and Across Sector

To understand how much of the observed decline in aggregate emissions is due to reductions in aggregate emissions intensity, we follow Levinson (2009) and decompose aggregate emissions into three channels: scale, composition, and aggregate technique. Letting m denote an industry or sector, total emissions, Z, in a given year, t, are given by:

$$Z_t = \sum_m Z_{mt} = \sum_m Q_{mt} \cdot E_{mt} = \underbrace{Q_t}_{\text{Scale}} \times \sum_m \underbrace{\Theta_{mt}}_{\text{Composition}} \cdot \underbrace{E_{mt}}_{\text{Aggregate Technique}} \tag{1}$$

where $\Theta_{mt} = Q_{mt}/Q_t$ represents the market share held by the sector and $E_{mt} = Z_{mt}/Q_{mt}$ represents the aggregate emissions intensity of the sector. The decomposition can also be expressed in vector notation for a given year: $Z = Q\Theta' E$, where Γ and E are both $m \times 1$ vectors capturing sector market share and corresponding emissions intensity. Totally differentiating this equation and dividing by Z, percent changes in total emissions can be expressed as the sum of percent changes in country size or total output ("scale"), changes in the relative market shares of cleaner and dirtier sectors ("composition"), and changes in aggregate emissions

[11] This contrasts with work by Greenstone (2003) that finds no evidence of increases in emissions to other media in response to the Clean Air Act Amendments.

intensity ("aggregate technique"):

$$\frac{\mathrm{d}Z}{Z} = \frac{\mathrm{d}Q}{Q} + \frac{\mathrm{d}\Theta}{\Theta} + \frac{\mathrm{d}E}{E} \tag{2}$$

Note that if we were to only consider a single sector, Γ_{mt} would equal one in each year, and there would be no composition effect. In this case, changes in aggregate emissions over time would be driven solely by the scale and aggregate technique effects.

2.2 Within and Across Establishment

To understand how establishment-level adjustment can affect aggregate emissions, working through the aggregate technique effect, we extend the decomposition in (1). We add an additional subscript, i, to denote individual establishments. We represent establishment-level variables with lower-case letters and aggregate variables with upper-case letters.

Emissions in a given sector are the sum of the emissions from each establishment in that sector, $Z_{mt} = \sum_{i \in m} z_{imt}$, we can further decompose the aggregate emissions intensity component from (1) as follows:

$$E_{mt} = \underbrace{\sum_i \underbrace{\theta_{imt}}_{\text{Reallocation}} \cdot \underbrace{e_{imt}}_{\text{Pure Technique}}}_{\text{Aggregate Technique}} \tag{3}$$

where $\theta_{imt} = q_{imt}/Q_{mt}$ represents an establishment's share of sector output, and $e_{imt} = z_{imt}/q_{imt}$ represents a establishment's emissions intensity. The "Reallocation" effect captures changes in aggregate emissions intensity that occur as establishments enter and exit and resources are systematically reallocated between establishments–an across-establishment effect. The "Within Establishment Technique" effect captures adjustment in establishment's production techniques–a within-establishment effect. Written in this way, aggregate emissions intensity in each sector and year is an output-share weighted-average of each establishment's emissions intensity. Since the establishment-level analysis is inherently a sub-sector analysis, we drop the sector subscripts through the remainder of the discussion.

To understand the relative importance of reallocation, in its various forms, and within-establishment production technique adjustment in driving aggregate emissions, our empirical approach will decompose the percentage change in aggregate emissions intensity over time (from $t = 1$ to 2): $\Delta E/E = (E_2 - E_1)/\overline{E}$, where the weight $\overline{E}$ is included to express the change in scale-independent percent-change terms.[12] To accomplish this additional inves-

[12] $\overline{E} = (E_1 + E_2)/2$

tigation, we modify an approach suggested by Melitz and Polanec (2015),[13] to decompose changes in aggregate emissions intensity into three channels–the changes due to survivors, entrants, and exiters–and then further decompose the surviving-establishment channel into changes due to across-establishment reallocation and within-establishment adjustment.

Let $\Theta_{Gt} = \sum_{i \in G} \theta_{it}$ represent the aggregate market share of a group, G, of establishments, where the G represents survivors (S), exiters (X), or entrants (N). Then define $E_{Gt} = \sum_{i \in G} (\theta_{it}/\Theta_{Gt})\, e_{it}$ as the group's emissions intensity. Aggregate emissions intensity in periods 1 and 2 can now be expressed as a function of the aggregate output share and aggregate emissions intensity of the three groups of establishments (survivors, entrants, and exiters):

$$\begin{aligned} E_1 &= \Theta_{S1}E_{S1} + \Theta_{X1}E_{X1} = E_{S1} + \Theta_{X1}(E_{X1} - E_{S1}) \\ E_2 &= \Theta_{S2}E_{S2} + \Theta_{N2}E_{N2} = E_{S2} + \Theta_{N2}(E_{N2} - E_{S2}) \end{aligned} \tag{4}$$

The final step of the decomposition builds on Olley and Pakes (1996) using the alternative decomposition of aggregate emissions intensity:

$$\begin{aligned} E_t &= \bar{e}_t + \sum_i (\theta_{it} - \bar{\theta}_t)(e_{it} - \bar{e}_t) \\ &= \bar{e}_t + \mathrm{cov}(\theta_t,\, e_t) \end{aligned} \tag{5}$$

where $\bar{e}_t$ is the unweighted average establishment emissions intensity and $\bar{\theta}_t$ is average market share.[14] In this way, changes in aggregate emissions intensity can be expressed as the sum of the change in unweighted average emissions intensity, $\Delta\overline{ei}$–this can be thought of as a within-establishment effect that is common to all establishments–and the change in the covariance (between market share and emissions intensity), Δcov–which can be thought of as a cross-establishment reallocation effect. As discussed by Melitz and Polanec (2015), expressing the results in scale-independent terms when decomposing data measured in levels, as we do, will also require a scale-independent covariance measure that will also be invariant to proportional changes in emissions intensity. We follow their lead in defining such a measure as $\widetilde{\mathrm{cov}} = \mathrm{cov}(\theta,\, e/E) = \mathrm{cov}(\theta,\, e)/E$. Thus, $\widetilde{\mathrm{cov}}$ represents the share of aggregate emissions intensity, E, driven by the correlation between market share and emissions intensity, a cross-establishment share, and the remaining share, $\bar{e}/E$, captures the share due to average emissions intensity, independent of its correlation with market shares.

[13] Melitz and Polanec (2015) use a similar approach to decompose the channels underlying changes in aggregate productivity over time.

[14] Melitz and Polanec (2015) note that the use of the covariance operator, which would typically be multiplied by $1/n_t$, is a slight abuse of notation, but, because θ_{nt} are shares, the equation basically already incorporate this division.

Combining the equations in (4) with the decomposition in (5), the change in aggregate emissions intensity is thus given by:

$$\begin{aligned}\frac{\Delta E}{\overline{E}} &= \frac{E_{S2}-E_{S1}}{\overline{E}} + \frac{\Theta_{N2}(E_{N2}-E_{S2})}{\overline{E}} + \frac{\Theta_{X1}(E_{S1}-E_{X1})}{\overline{E}} \\ &= \underbrace{\frac{\Delta \overline{e}_S}{(1-\overline{\widetilde{\text{cov}}}_S)\overline{E}}}_{\text{Pure Technique}} + \underbrace{\frac{\Delta \text{cov}_S}{(1-\overline{\widetilde{\text{cov}}}_S)}\frac{\overline{E}_S}{\overline{E}} + \frac{\Theta_{N2}(E_{N2}-E_{S2})}{\overline{E}} + \frac{\Theta_{X1}(E_{S1}-E_{X1})}{\overline{E}}}_{\text{Reallocation}} \end{aligned} \tag{6}$$

where $\overline{E}_S = (E_{S2}+E_{S1})/2$ and $\overline{\widetilde{\text{cov}}}_S = \left(\overline{\widetilde{\text{cov}}}_{S2} + \overline{\widetilde{\text{cov}}}_{S1}\right)/2$ represent the time average over periods 1 and 2.

The first line decomposes the percent change in aggregate emissions intensity into the share due to survivors, entrants, and exiting establishments using (4). The second line further decomposes the change due to survivors into the change in the distribution of emissions intensity (which can be thought of as a within-establishment adjustment in production techniques that is common to all surviving establishments) and the change due to market share reallocation between cleaner and dirtier establishments (an across-establishment reallocation of resources among survivors).[15] This decomposition has several distinct advantages.

First, the decomposition in the second line cleanly separates changes in aggregate emissions intensity into four possible channels: within-establishment adjustment to production techniques and emissions intensity, reallocation among surviving establishments, reallocation to entrants, and reallocation away from exiting establishments. Building on the aggregate decomposition in equation (2) commonly used in related literature, our additional establishment-level decomposition results can be easily linked to changes in aggregate emissions.

Second, as noted by Melitz and Polanec (2015), this approach leverages the cross-sectional nature of the Olley and Pakes approach. Thus the decomposition of the emissions changes over time into three groups–survivors, exiting, and entering establishments–need not use the same reference emissions intensity value for each group. The decomposition is only constrained so that the sum of the three changes sum to the actual total change. Other decomposition approaches used to examine productivity, for example Foster et al. (2001), include a fixed reference group–either establishments from a single period, or an average from multiple periods–and thus tend to miss trends in productivity, which Melitz and Polanec (2015) argues introduces bias, tending to understate the relative contribution of survivors.

Finally, as formulated, the three channels have an intuitive interpretation. For example, the change due to survivors is the change in aggregate emissions intensity that would

[15] The same Olley and Pakes decomposition could be extended to entering and exiting establishments as well.

have occurred if there were no entry and exit. Then, using surviving establishments as a benchmark, the change due to entry, $\Theta_{N2}(E_{N2} - E_{S2})$, is the change in aggregate emissions intensity that would occur from adding or subtracting the entrants. Thus, entering establishments will contribute to a decline in aggregate emissions intensity if they have lower aggregate emissions intensity than survivors in period two, $E_{N2} < E_{S2}$. Similarly, exiting establishments will contribute to a decline in aggregate emissions intensity if they have a higher aggregate emissions intensity than the surviving establishments they leave behind, $EI_{X1} > EI_{S1}$.

3 Empirical Strategy

Our empirical analysis of the drivers of the reduction in toxic pollution emissions from the U.S. manufacturing sector proceeds in two steps. First, we decompose the change in pollution emissions into the traditional scale, composition and technique effects. We then further decompose the technique effect into the four channels described above: reallocation among entrants, exiters, and survivors, and the within establishment technique effect. In this section we briefly describe the data and the empirical approach we employ.

3.1 Data

Our approach to measuring the drivers of the fall in toxic pollution emissions requires data on output and emissions at the establishment level. For pollution emissions we employ the EPA's Risk Screening Environmental Indicators database (RSEI) based on emissions reported in the Toxic Release Inventory (TRI).[16] To measure establishment output we use the National Establishment Time Series (NETS). In this subsection we briefly describe both data sets and the process used to link them. The merged data set is an unbalanced panel of establishment-year observations on U.S. manufacturing plants over the years 1990-1999.[17] For each establishment we observe output, employment and industry. We are able to match these establishments to their reported TRI emissions and are able to observe quantities of pollution over 650 different chemicals by media of release (air, water, and various offsite transfers). A version of this data set was employed to assess the relationship between environmental performance and export orientation in Holladay (2016).

[16] In general, we refer to TRI emissions, which are the direct source of the emissions. The RSEI database makes use of these reported TRI emissions to additionally evaluate potential hazards to human health based on chemical toxicity estimates.

[17] We focus on manufacturing establishments since they represent the vast majority of emissions reported in the RSEI database, and because we can make use of the NBER-CES manufacturing-sectors price indices to calculate a measure of real output from the sales data we observe in the NETS database.

The National Establishment Time Series (NETS) is a proprietary database compiled from Dunn and Bradstreet data on the creditworthiness of establishments. It claims to provide data on the universe of U.S. establishments. The NETS distributes the Dun and Bradstreet data to researchers and companies for market research. A number of papers in the economics literature have used the NETS.[18] Most notably, Neumark et al. (2011) finds that the data in the NETS is comparable in quality to other public and private data sets, including the US Census.[19] Dunn and Bradstreet collect data on output, employment, location and ownership and use it to create a credit rating for establishments. The NETS includes data on the dollar value of the establishment's sales, employment and the Dunn and Bradstreet credit rating. It also records the company's primary industry and product (8-digit SIC) and up to five secondary products.

The TRI tracks the disposal of hundreds of different toxic chemicals regulated under the Emergency Planning and Community Right-to-Know Act (EPCRA). The TRI records the annual amount of toxic chemicals disposed by media of disposal, facility and chemical. Facilities are required to report their toxic emissions if they: (1) have more than ten employees, (2) "produce, process or otherwise use" more than a threshold level of any single regulated chemical,[20] and (3) are in the manufacturing sector, or a handful of other related sectors.[21] The TRI has been widely used in the economics literature, but they are an imperfect measure of toxic pollution emissions. de Marchi and Hamilton (2006) and Koehler and Spengler (2007) demonstrate evidence of under-reporting in the TRI.[22] TRI emissions are reported by the facilities themselves, but EPA is authorized to ensure compliance and brings a number of cases each year against polluters who misreport.[23] We are not aware of any evidence that under-reporting or enforcement activities changed during our sample period. The RSEI distributes the TRI data along with toxicity weights for each of the reported chemicals, which

[18] See Levine et al. (2012), Cui et al. (2016) and Cherniwchan (2017), for examples. Brinkman et al. (2015) uses a different version of the Dun and Bradstreet data on which the NETS is based.

[19] Though, as Haltiwanger et al. (2013) observe, the NETS database includes firms with and without employees and does not appear to have complete coverage of both types of firms. The annual total number of firms and establishments in the NETS is generally fewer than is documented in the US Census Longitudinal Database (LDB).

[20] The most common threshold is 10,000 pounds but more toxic chemicals have lower thresholds ranging down to 0.1-grams for dioxin. The reporting threshold is defined by use, not emissions, so many facilities report emissions far below the threshold.

[21] In this paper we restrict our attention to the manufacturing sector.

[22] de Marchi and Hamilton (2006) finds that reported reductions in TRI emissions are not always matched by reductions in pollution concentrations at air monitors and those reported emissions are not consistent with Benford's Law for two pollutants. Koehler and Spengler (2007) conducts a case study in aluminum industry and finds evidence that emissions of polycyclic aromatic hydrocarbons (PAH) are under-reported after the introduction of new regulations.

[23] TRI compliance history for facilities is available from EPA's Enforcement and Compliance History Online (ECHO) database.

allows us to evaluate changes in emissions by quantity or toxicity.

Much of the literature that assesses the fall in pollution levels from manufacturing sector uses the National Emissions Inventory (NEI) as the source of pollution data. The NEI claims to be a comprehensive measure of point source air polluters without any of the reporting requirements that affect the TRI. We believe the TRI is more appropriate for our analysis because it reports annually while the NEI is triennial. Annual data allows us to better evaluate the impacts of entry and exit on aggregate pollution emissions. The TRI also reports emissions across different media, which is helpful in assessing whether the fall in air pollution masks shifts in pollution to water or landfills. Finally, the TRI allows us to confirm the results of the existing literature using a different data set covering pollutants regulated in a very different manner.[24]

The set of chemicals that are reported in the TRI has changed over time. In this analysis we restrict our focus to the emissions of "core" TRI chemicals that have been regulated since the TRI began in 1988. We aggregate all reported toxic chemical emissions from an establishment in a year and use that as our measure of pollution emissions.[25] We then match establishments from the TRI to the NETS. The matching procedure is described in Holladay (2016) and detailed in the appendix below. The final matched data set contains 92,210 establishment-year observations, comprised of 16,226 manufacturing establishments reporting emissions over the sample period.

Because we focus on the contribution of establishment entry and exit to the cleanup in manufacturing, accurately measuring establishment churn is particularly important. Importantly, the NETS data tracks when establishments enter and exit and provides a link to their parent firm.[26] The sample contains a number of multi-establishment firms and we define

[24]In addition, according to the EPA's report "Factors to Consider When Using Toxic Release Inventory Data" (2015), "79 percent of hazardous air pollutant releases found in the 2002 NEI data set were also documented in TRI" (p.11).

[25]The EPA has published a concordance that links numerous TRI chemicals to a corresponding criteria air pollutant (Environmental Protection Agency, 2013). Using this concordance we classified establishment emissions as Volatile Organic Compounds (VOCs), as well as SO2 and particulate matter (PM-10). Grouped in this way, VOCs make up the largest single group in our data and the decomposition results are quite similar to the full sample results we report here. The PM-10 and SO2 observations make up a smaller share (17 percent and 4 percent, respectively) and the results are noisier but also similar. The results for VOCs, PM-10, and SO2 all show a positive scale plus composition effect and reductions in emissions due to declines in aggregate emissions intensity. VOC and PM-10 results exhibit declines due to both the technique and reallocation channels, while SO2 exhibits a small increase due to the reallocation channel. Results are available upon request.

[26]The NETS includes "firstyear" and "lastyear" fields identifying the first and last year that an establishment is active, respectively The NETS links establishments to their parent firm headquarters via an "HQDuns" field. The Dunn and Bradstreet database uniquely identifies establishments and headquarters using a "DUNSnumber" and "HQDuns", respectively. All establishments in the same firm report the same "HQDuns" number.

entry and exit as firm-level outcomes. Thus, we say an establishment "exits" if that establishment's last active year coincides with the year the last remaining establishment reporting that headquarters DUNS value ceases to be active and exits the sample. This approach will capture the environmental effects of any establishment exit that does not coincide with the firm's exit in the "survivors" group, identified in equation (6).[27]

3.2 Trends in toxic emissions

Toxic pollution emissions have been falling consistently since they began being tracked by the TRI in 1988. Figure 1 shows the trends in real output and toxic pollution emissions, represented by the sum of TRI emissions reported by manufacturing establishments in our matched sample, from 1990-1999. Output rose 13 percent during that time while emissions fell by 52 percent.[28] The trend has continued, with aggregate TRI emissions falling by a third between 2000 and 2015. The fall in aggregate TRI emissions is reflected in falling emissions of the individual chemicals that are reported. In addition, hazard scores, which weight the pounds of emissions by the toxicity of the chemical emitted, have fallen by roughly a quarter, representing a significant reduction in the risk to human health. This suggests that changes in the set of chemicals emitted has not undone the environmental benefit of reduced emissions levels. Emissions of the five most toxic chemicals are down by a third.[29] This is roughly consistent with the trend in criteria air pollutants regulated under the Clean Air Act. The National Emissions Inventory reports that emissions of criteria pollutants fell between 23 percent and 50 percent between 1990 and 2000 depending on the pollutant.[30]

[27] The environmental effects of these establishment closures would then be further decomposed according to the second line of equation (6).

[28] Real output is calculated using sales data from establishments in the NETS sample and deflated by industry-specific price indices published in the NBER-CES database described in Becker et al. (2013). See the appendix for more details on data set construction.

[29] The five most toxic chemicals based on the toxicity weights posted in the RSEI are Bis(chlorophyll) ether, Thorium dioxide, Propyleneimine, Asbestos (friable), and Hydrazine. The reductions in emissions range from 100 percent for Thorium dioxide to 16 percent for Asbestos.

[30] Source: author's calculations using EPA's National Emission Inventory Air Pollutant Emissions Trends data (https://www.epa.gov/air-emissions-inventories/air-pollutant-emissions-trends-data), version 1970-2016, for the category "Industrial and other processes".

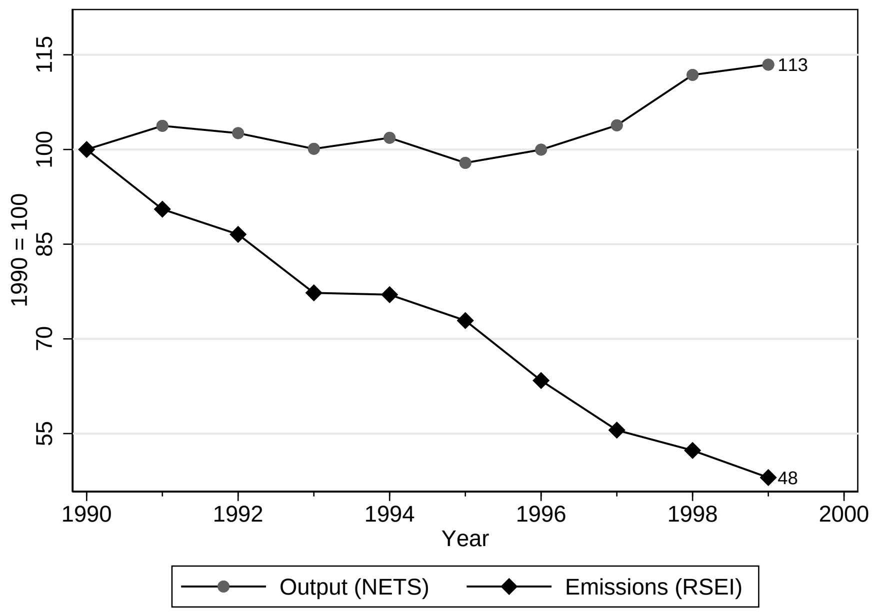

Figure 1: U.S. Manufacturing Real Output and Emissions

Note: Emissions is an aggregate measure of the pounds of emissions of TRI chemicals released by all establishments as reported by EPA's RSEI database. Real output is calculated using sales data from establishments in the NETS sample and weighted by SIC4 price indices published in the NBER-CES database described in Becker et al. (2013).

3.3 Decomposing the trend in toxic emissions

To better understand the drivers of the improved environmental performance in the U.S. manufacturing sector, we follow the literature by first decomposing the annual scale, composition, and aggregate technique channels in our matched NETS and TRI data. The percent change in total emissions can be expressed as the sum of the percent changes in each of the scale, composition, and aggregate technique effects, as shown in equation (2). We then implement our further decomposition of the percent changes in the aggregate technique effect into the within-establishment technique and across-establishment reallocation effects, given in equation (6). The results are summarized in Figure 2, where values are scaled to equal 100 in 1990.

First, the scale effect, captured by changes in real output, rises 13 percent, shown by line (1). This suggests that holding the composition and emissions intensity of manufacturing in our sample constant we would have expected emissions of toxic chemicals to increase by 13 percent. This outcome is somewhat lower than the levels reported in the literature from estimates of changes in criteria air pollutants.[31] To understand the role of the composition

[31] For example, the aggregate NBER-CES data reports a 23 percent increase in real manufacturing output from 1990-1999. The scale effect documented using our establishment-level NETS sample is approximately ten percentage points lower.

effect, we first measure emissions intensity by SIC4 industry in 1990.[32] We then calculate the change in emissions implied by shifts in the cross-sector output shares by multiplying the fixed 1990 emissions intensity measure by sector-output in each year. This identifies the combined scale and composition effect, holding sector-level aggregate emissions intensities fixed at their 1990 levels. The additional change due to the composition effect shown in Figure 2 implies an additional increase in emissions, as shown by line (2). If the emissions intensity of U.S. manufacturing output had remained constant over our sample period, emissions would have increased by 10 percentage points. This suggests that there has been a shift in production towards relatively dirty industries.[33] Adding in the within-establishment technique effect, emissions would have declined to 65 percent of their 1990 level, shown by line (3). The within establishment reduction in toxic emissions intensity accounts for a decline in emissions of 58 percent[34]–more than one hundred percent of the observed fall in toxic pollution emissions from the U.S. manufacturing industry over our sample period.[35]

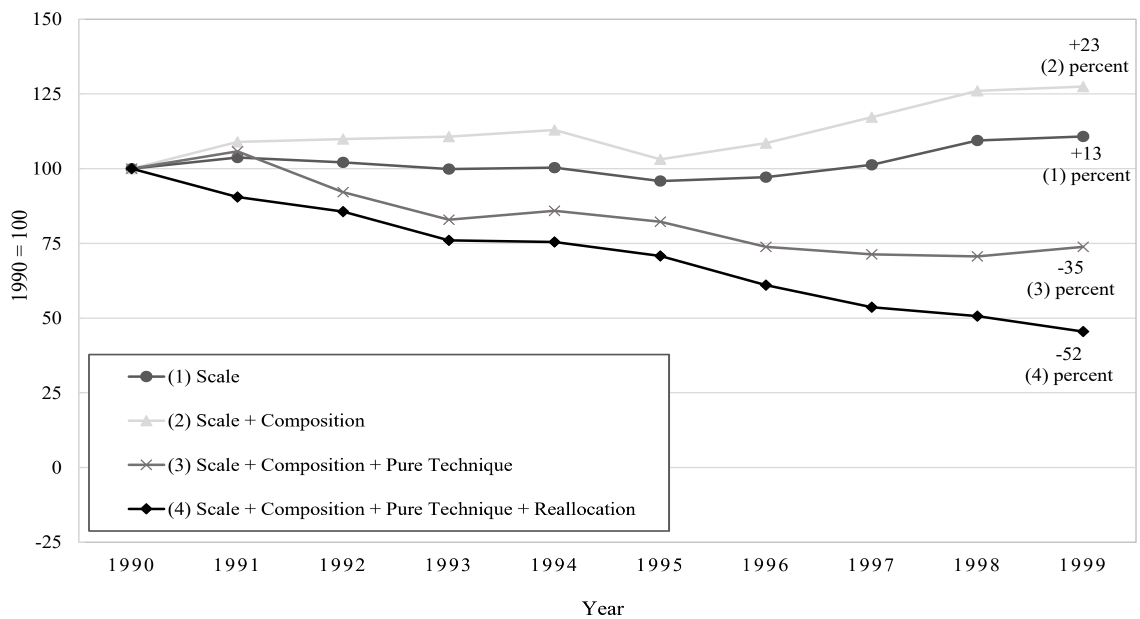

Figure 2: Toxic Emissions from U.S. Manufacturing

Note: This figure plots the percent-change in emissions, collected in the Toxic Release Inventory, decomposed into the 4-channels defined above. The fourth line represents the sum of all four effects and reports the percent-change in overall emissions relative to the base year of 1990.

Finally, line (4) reports the change in total toxic emissions by adding in the the realloca-

[32]We measure emissions intensity by sector in 1990, as shown in (1), using data from our sample.

[33]This outcome is in contrast to the composition shift towards cleaner sectors identified by Levinson (2009) and Shapiro and Walker (2016) using NEI data.

[34]-35 percent minus 23 percent.

[35]Similarly, Levinson (2015) finds that, for some criteria pollutants, the aggregate technique effect accounts for more than 100 percent of the observed emissions reduction.

tion effect, changes in emissions as establishments enter and exit and as output share shifts among surviving establishments. Cross-establishment reallocation represents an additional important cleanup channel, driving down emissions by an additional 17 percent.[36] The reallocation effect plus the within-establishment technique effect yields the changes due to the aggregate technique effect, and taken altogether the sum of the four channels (scale, composition, technique, and reallocation) yields the actual decline in emissions over the sample period. From 1990-1999, total toxic emissions declined by 52 percent, with the aggregate technique effect accounting for 75 percent of the implied drop. Our combined results are generally consistent with existing literature. For example, Levinson (2009) finds that from between 1987-2001 emissions declined by 27 percent, and the aggregate technique effect accounted for a drop in emissions of 39 percent–thus, accounting for more than 100 percent of the observed reduction in emissions. Shapiro and Walker (2016) present evidence that the aggregate technique effect accounts for at least 63 percent and as much as 104 percent of the reduction in criteria air pollution emissions, depending on the pollutant under consideration. Our analysis provides the additional insight that these aggregate reductions in emissions intensity are not driven by within-establishment technology upgrading and abatement activities, but are also driven by reallocation of market activity towards cleaner firms and establishments.

3.4 Decomposing the technique effect

As described in equation (6) above, we develop a method to decompose the aggregate technique effect into four separate components: reallocation among surviving, entering establishments, exiting establishments, and the within-establishment technique effect. To better understand the factors driving the reallocation and technique channels presented in Figure 2, we present these additional decomposition results in two tables. The first set of results based on equation (6), are presented in Table 1, which reports the change in aggregate emissions intensity due to survivors (which are additionally broken into the within-establishment technique and between-establishment reallocation channels), entrants, and exiting establishments, and their sum from 1991 to 1999.[37] The decline documented from year to year in the total corresponds to the annual changes in aggregate emissions intensity that can be seen from year to year in the decline of line (4) relative to line (2) shown in Figure 2. The annual changes reported are in scale-independent terms and can be interpreted as percentage

[36] -52 percent minus -35 percent yields the additional reduction of 17 percent due to the reallocation effect.

[37] Note that the total change due to surviving establishments is equal to the sum of the within and between establishment channels, and the total change due to all establishments is the sum of the changes due to survivors, entrants, and exiters.

point changes. Table 2, reports the within-establishment technique effect and the reallocation effect, which is the sum of the reallocation between surviving establishments and those entering or exiting in each period.[38]

Examining the contribution of each channel in Table 1, several trends emerge. First, the cleanup among surviving establishments is the most important factor driving observed declines in aggregate emissions intensity and the sign correlates with the aggregate change every year except 1995. Among these survivors, cleanup within establishments as well as reallocation of productive activity towards cleaner establishments are both important channels driving down emissions. However, both channels do not contribute to reductions in every year, for example, in 1992 reallocation of market share toward dirtier establishments generated a drag on the cleanup among survivors that would otherwise have occurred if driven solely by the within-establishment technique effect.

Second, entering establishments are very similar to the average establishments they are entering alongside and do not exhibit a large effect, in either direction, on the aggregate emissions intensity. Finally, exiting establishments and firms are dirtier than the surviving establishments they leave behind, on average,[39] though, again, this is not a consistent phenomena over time. In some years, the exiting establishments are cleaner than the firms they leave behind, for example in 1991, 1993, etc. In these years, this exit is a drag on the clean up of toxic pollution emissions from manufacturing during our sample period. In other years, when exiting firms are dirtier, they contribute to additional cleanup. Overall, the within-establishment technique effects and and the reduction in emissions due to cross-establishment reallocation is driven primarily by market activity among surviving firms. Our results demonstrate that, while important, entry and exit exhibit a relatively smaller effect on emissions outcomes. To get a sense of the relative importance of each channel we computed the mean of the absolute value of each column. Survivors average 7.5, entrants average 0.2, and exiters average 1.8. Thus, as a share, survivors account for 78 percent, entrants for 3 percent, and exiters for 19 percent of the observed changes in aggregate emissions intensity over the sample period. Together, entry and exit appear to drive roughly 22 percent of the changes in aggregate emissions intensity, with the remainder driven by surviving establishment activity.

Table 2 reformulates the aggregate emissions intensity decomposition, to give the within-

[38] Differences in the Total column in each table are only due to rounding differences.

[39] Recall from equation (3) that exiting establishments will contribute to a decline in aggregate emissions intensity if exiting establishments have a higher emissions intensity than the establishments they leave behind. The positive values in the exiting establishments column are tending to raise aggregate emissions intensity as exiting establishments, on average, have a lower emissions intensity than incumbents they leave behind in those years. Due to the sign changes over the period, the average impact of exiting establishments is an emissions reduction of 0.2 percent.

Table 1: Change in Aggregate Emissions Intensity by Channel

Year	Surviving Firms Total	Surviving Firms Within	Surviving Firms Between	Entering Firms	Exiting Firms	Total (All Firms)
1991	-17.2	-8.9	-8.4	-0.1	4.1	-13.2
1992	-2.2	-9.6	7.4	0.0	-0.6	-2.8
1993	-9.4	-8.5	-0.9	1.1	1.6	-6.8
1994	-2.7	1.0	-3.7	0.3	0.4	-2.0
1995	1.5	4.8	-3.3	0.0	-1.6	-0.1
1996	-9.8	-11.1	1.3	0.3	-2.1	-11.6
1997	-12.7	-11.2	-1.5	0.0	0.9	-11.7
1998	-7.7	-6.9	-0.9	-0.2	-3.3	-11.2
1999	-4.7	1.4	-6.1	0.2	-1.4	-5.9

Note: Underlying data are establishment-year observations. Entry and Exit are identified at the firm level. Establishments are considered to have exited if their parent firm also exits; similarly for entry.

surviving-establishments technique effect and the reallocation effect in each year, by adding together the three components of the reallocation effect, given in equation (6). The second column reports the within-establishment percentage-point change due to survivors, which is the technique effect.[40] The third column sums the change due to reallocation across surviving establishments, exiting establishments, and entering establishments. The final columns reports the sum of the two, which are the changes in aggregate emissions intensity, or the aggregate technique effect in each year.

During much of the sample period the two channels are of similar magnitude and the net result of churning was to move resources towards relatively less emissions-intense establishments, two exceptions being 1992-1993. On average, the within-establishment technique channel has been more important than reallocation across establishments in driving observed changes in aggregate emissions intensity, but the reallocation of resources across establishments has been an important factor in every year, and the average trend has been to reallocate activity towards cleaner establishments.

4 TRI Emissions Media Categories

The existing literature analyzing the fall in manufacturing pollution emissions has focused largely on emissions of the six pollutants regulated under the Clean Air Act and its amendments. One of the driving questions of this literature has been the extent to which the

[40]This information is also reported in the second column of results in Table 1.

Table 2: Reallocation and Within-Establishment Technique Effects

Year	Within Establishment (Technique)	Across Establishment (Reallocation)	Total (Aggregate Technique)
1991	-8.9	-4.3	-13.2
1992	-9.6	6.8	-2.8
1993	-8.5	1.7	-6.8
1994	1.0	-3.1	-2.1
1995	4.8	-4.9	-0.1
1996	-11.1	-0.5	-11.6
1997	-11.2	-0.5	-11.7
1998	-6.9	-4.3	-11.2
1999	1.4	-7.3	-5.9

Note: Change in aggregate emission intensity decomposed into the two channels that comprise the aggregate technique effect. Units are the percent of 1990 toxic release inventory emissions in our sample.

Clean Air Act itself has caused the fall in pollution emissions from U.S. manufacturers. One advantage to employing the RSEI database, based on TRI data, is that it allows us to examine the evolution of emissions of a different set of pollutants into a number of different media. The toxic chemicals covered under the TRI are regulated under a different regulatory framework[41] and the TRI requires detailed information on how chemicals are disposed. Each establishment must report the quantity of each chemical released via each of twenty-one different disposal media ranging from onsite recycling to direct water discharge.[42]

We implement the empirical strategy described above for the two major release categories accounted for in the TRI database: air and water.[43] The vast majority of the emissions documented in the TRI database fall into one of these media of release. For ease of exposition we present the results graphically for each disposal channel using a figure analogous to Figure

[41]Toxic chemicals are primarily regulated under the Toxic Substances Control Act (TSCA) and the Emergency Planning and Community Right-to-Know Act of 1986 which established the TRI. These regulations focus on information provision and risk management rather than the command and control style regulation in the Clean Air Act and its amendments.

[42]Major categories include air, water, on site land releases and offsite transfer for disposal. Each of these major categories can include multiple subcategories.

[43]The RSEI database makes use of six of the release-channels reported to the TRI to estimate the potential human health risks of different chemical releases: Fugitive Air, Stack Air, Direct to Water, Transfer to Publicly Owned Treatment Works (water), Offsite for Incineration, Offsite for Incineration (no fuel value). We have aggregated the two air channels and the two water channels. The two incineration channels account for 9 percent of RSEI emissions in our sample and their decomposition is very similar to air emissions. However, these release channels may include chemicals that were stored up over several years, and is thus not necessarily limited to chemicals used or produced in the reporting year. Results are available for any of the emissions media by request.

2.

Figure 3 presents the decomposition results for onsite releases to the air, which account for 78 percent of total TRI emissions and 71 percent of the observations in our sample. The results are similar to those presented in figure 2. Overall airborne emissions have fallen by 58 percent over the sample period, compared to a 52 percent fall in total TRI emissions. The scale effect accounts for an 11 percent increase, slightly less than the overall TRI because not all TRI reporters emit pollution to the air. The composition effect is also essentially unchanged from the aggregate TRI results. The technique effect accounts for a 69 percent reduction[44] and the reallocation effect drives emissions down by an additional 12 percent.

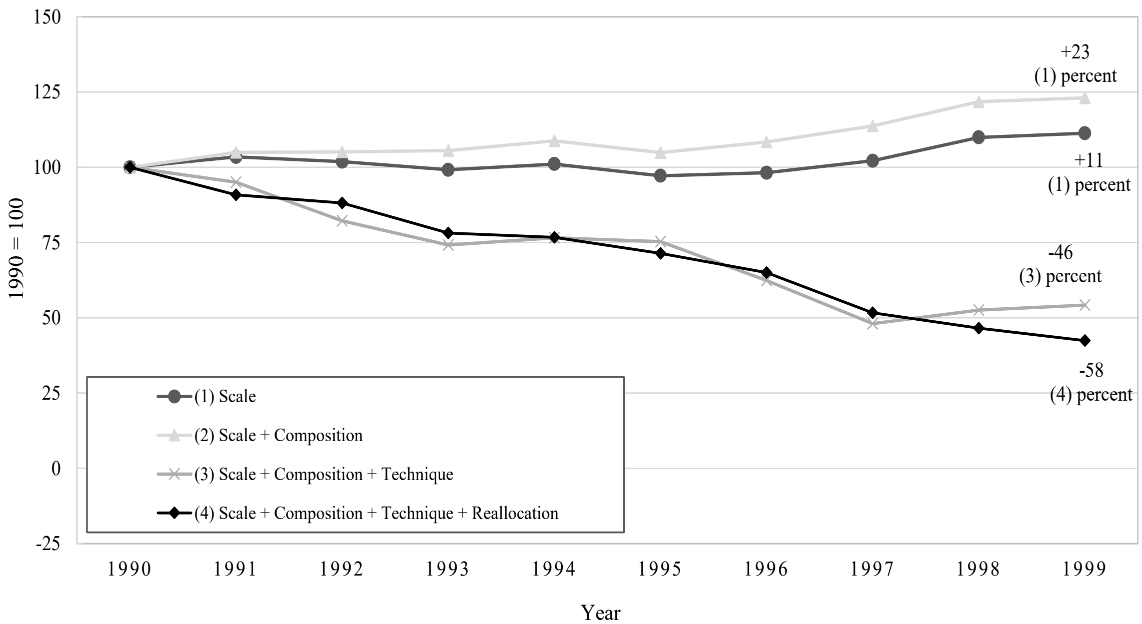

Figure 3: Emissions to Air from U.S. Manufacturing

Note: This figure plots the percent-change in emissions, collected in the Toxic Release Inventory that were released into the air, decomposed into 4-channels. The fourth line represents the sum of all four effects and reports the percent change in overall emissions relative to the base year of 1990.

We next analyze water releases, which account for 13 percent of total toxic emissions. In this category we aggregate onsite releases to surface water bodies and offsite transfers to publicly owned water treatment works. Offsite transfer to treatment works represents 92 percent of the releases to water, with direct releases accounting for the remaining 8 percent. Thus, changes in releases to water over the period are largely driven by changes in the transfer to publicly owned treatment works. Figure 4 decomposes these toxic emissions. Here the pattern is somewhat different than for airborne emissions. Overall water discharges fell by 35 percent, broadly consistent with Shapiro and Walker (2017) findings that broad measures of

[44] -46 percent minus 23 percent.

water pollution (such as fishability) improved during this time frame. While significant, the decline in water discharges of toxic waste is much smaller than the fall in overall toxic emissions. The scale effect is larger, implying that water polluter's output is growing faster than the average toxic polluter in our sample. Most notably, the within-establishment technique effect is much smaller for toxic releases to water. The within establishment technique effect contributes 25 percent towards the reduction in toxic water emissions.[45] The reallocation effect, however, is relatively much larger, accounting for an additional decline in emissions of 38 percent. As with the airborne emissions, reallocation of water discharges has been towards relatively cleaner establishments, which has aided in the cleanup of these releases.

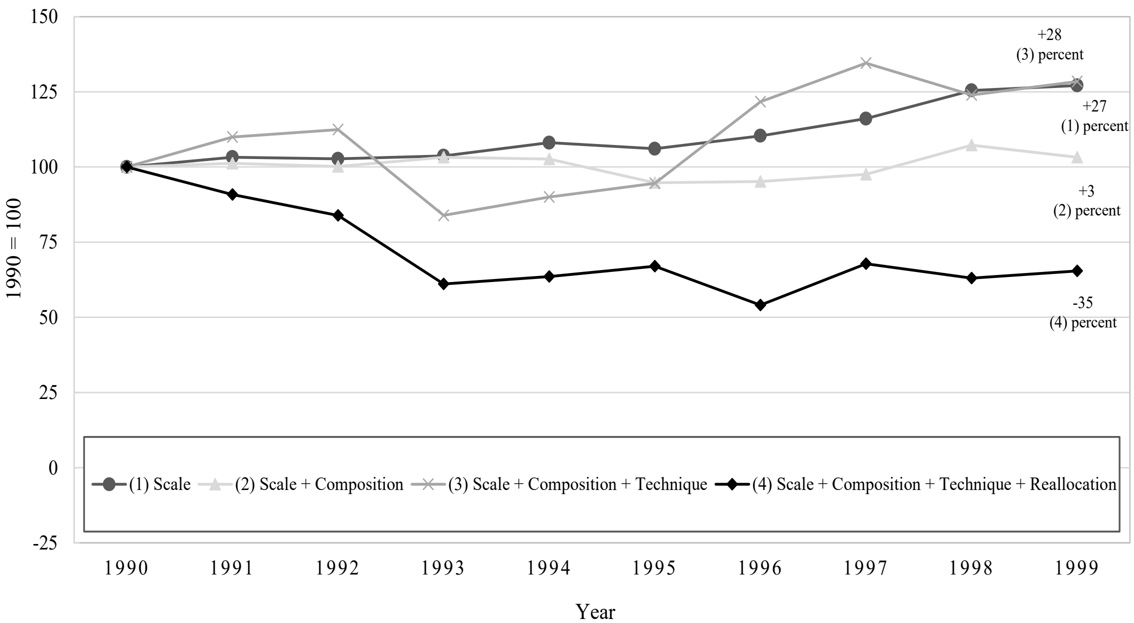

Figure 4: Emissions to Water and Treatment Works from U.S. Manufacturing
Note: This figure plots the percent-change in emissions, collected in the Toxic Release Inventory that were released into the water or transferred to a treatment works, decomposed into 4-channels. The fourth line represents the sum of all four effects and reports the percent change in overall emissions relative to the base year of 1990.

For each of the methods of disposal we analyze here, the total amount of toxic waste disposed has decreased. The composition effects have have been a drag on the reduction in emissions across the full sample period, as economic activity has tended to shift towards relatively dirty sectors. However, within-establishment improvements in techniques as well as across-establishment reallocation towards cleaners producers have both contributed substantially to observed declines in emissions across multiple release channels.

The decomposition results suggest that emissions into different media have evolved in different ways. Emissions have fallen further for air than water and the technique effect is

[45] +3 percent minus 28 percent.

relatively smaller, and reallocation effect relatively larger, for water discharges than the other disposal methods. While the literature has largely focused on the evolution of air pollution emissions from U.S. manufacturing, it appears that water emissions may have been reduced less than air. This is consistent with the hypothesis that environmental policy targeting air criteria pollutant emissions has had the ancillary benefit of reducing other types of air pollution emissions, but is also consistent with the concern that regulations on airborne emissions may drive firms to substitute to an alternate release channel like water.[46]

5 Robustness Checks

We explore the sensitivity of our results in two ways. First we analyze how weighting the quantity of toxic chemical by toxicity affects the decomposition result. We then consider the implications of different establishment weighting options.

The results presented thus far aggregate across different toxic chemicals by weight. The chemicals that are reported in the TRI vary tremendously in their toxicity and so the RSEI includes measures of toxicity by disposal method. The RSEI provides a "Hazard" metric that allows pollution emissions from each reporting establishment to also be ranked by toxicity.[47] As described in section 3.1, we collected establishment level hazard scores as a proxy for the toxicity of the emissions. We implement the decomposition described above to analyze the evolution in the hazard associated with pollution from U.S. manufacturing establishments across our sample period. The results are presented in Figure 5.

While the total weight of toxic emissions from U.S. manufacturing facilities reported in the TRI has fallen by 52 percent, toxicity weighted emissions are down by only 10 percent, as shown by line (4). This suggests that the fall in toxic emissions is masking a shift to more toxic chemicals and provides justification for decomposing the hazard scores as well. The scale effect is the same as the scale effect in unweighted emissions. This is expected since the scale effect, as shown in equation (1) holds emissions intensities within and across sectors fixed, and is essentially an aggregate output index. The composition effect is slightly larger when using toxicity weights, suggesting a stronger cross-sector shift of output towards sectors with relatively more toxic emissions. The technique effect is considerably larger for hazard (indicating a decline of 89 percent) while the reallocation effect is actually positive–indicating a drag on the cleanup as productive activity is being reallocated towards establishments producing or making use of more toxic chemicals.

[46] This idea is supported by recent work by Gibson (2016).

[47] Of the approximately six hundred chemicals in the TRI over four hundred have associated toxicity weights.

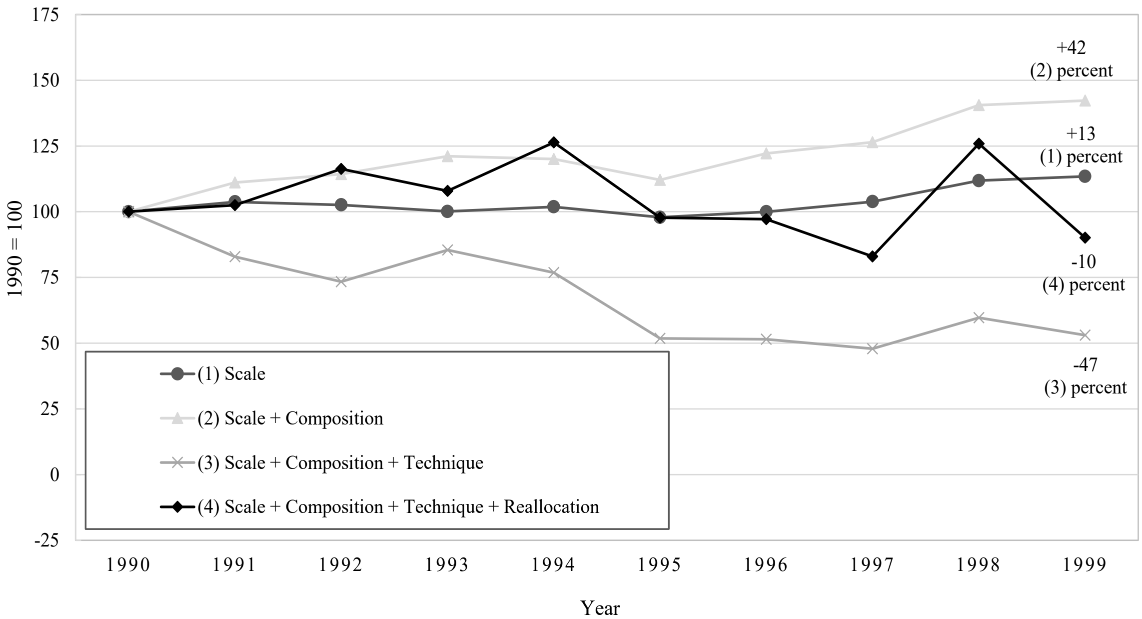

Figure 5: Hazard Score from U.S. Manufacturing Emissions

Note: This figure plots the decomposition of the EPA Hazard Score, a risk weighted measure of pollution emissions. The fourth line represents the sum of all four effects and reports the change in overall emissions relative to the base year of 1990.

As with pollution emissions we further decompose the reallocation effect into the contribution of surviving, entering, and exiting establishments, in Table 3. There is significant variation in the yearly estimates, but the general trend is for surviving establishments to become cleaner. On average, reallocation of output has been towards more toxic polluters, though also with substantial annual variation. Exiting establishments tend to have lower hazard scores than their surviving competitors, meaning their exit increased industry emissions intensity. Entering establishments tend to have relatively low hazard scores compared to incumbents, but there are a few years with notable exceptions.

The second concern to address is the choice of establishment weights. First, as discussed by Melitz and Polanec (2015), there is no theoretical reason to prefer one establishment-weighting scheme over another when focusing on changes in aggregate productivity. The same is true in for an analysis of aggregate emissions intensity, but with one important caveat. We are not concerned, solely, with the channels driving changes in aggregate emissions intensity. As laid out in equation (1), aggregate emissions intensity is just one component driving aggregate emissions, which is our ultimate concern. To link our decomposition of aggregate emissions intensity back to this original decomposition requires that we maintain the same output-weighting approach throughout. Laying this consideration aside, though, there is no other reason to prefer output shares as weights in an analysis of aggregate emissions intensity, in isolation.

Table 3: Change in Aggregate Hazard Intensity by Channel

	Surviving Firms			Entering Firms	Exiting Firms	Total (All Firms)
Year	Total	Within	Between			
1991	-2.0	-20.8	18.8	-0.1	0.8	-1.3
1992	10.8	-8.4	19.2	0.2	4.0	15.0
1993	-8.7	14.5	-23.3	0.4	2.4	-5.9
1994	2.5	-10.3	12.8	11.8	2.4	16.7
1995	-19.6	-21.1	1.5	0.0	-5.1	-24.7
1996	-4.5	-2.4	-2.1	-0.3	2.1	-2.7
1997	-14.9	-7.4	-7.5	-0.2	-2.9	-18.0
1998	51.9	3.8	48.1	-0.5	-16.6	34.9
1999	-37.5	-8.3	-29.3	-0.2	0.4	-37.4

Note: Hazard, not Pounds

Further, as documented by Foster et al. (2001), the choice of different weighting approaches is known to generate different conclusions regarding the channels driving employment, productivity and various other outcomes studied in related literature. It would not be surprising if different choices of weights generated different conclusions in an analysis of aggregate emissions intensity. To explore this possibility we compare the trend in weighted-aggregate emissions intensity using four different establishment-level weights–nominal sales, employment, emissions, and emissions intensity–and compare them to our baseline aggregate result using output.

Aggregate emissions intensity is identical, whether calculated as the sum of establishment emissions in a year, divided by the sum of real establishment output, or whether calculated as the sum of establishment emissions intensity weighted by establishment real-output-share.[48] This equivalence provides our baseline result which we compare to four other weighting options: nominal sales, employment, emissions, and emissions intensity, in Figure 6.

Using real output shares gives more weight to larger establishments, in terms of output. Using employment shares, instead, gives more weight to larger employers. Examining the panels in the first row of of Figure (6), the differences between real and nominal sales weights are quite small, while the second panel suggests that, in contrast to larger producers, larger employers may tend to be a littler more emissions intense, with their emissions intensity declining more slowly, on average, over the panel. This suggests that the decline in aggregate emissions through the aggregate emissions intensity channel is driven by larger and more productive establishments, and implies that nominal sales and employment data do not

[48]For all polluting establishments in each year, $E_t = \sum_i z_{it} / \sum_i q_{it} = \sum_i (q_{it}/Q_t) \cdot (z_{it}/q_{it})$

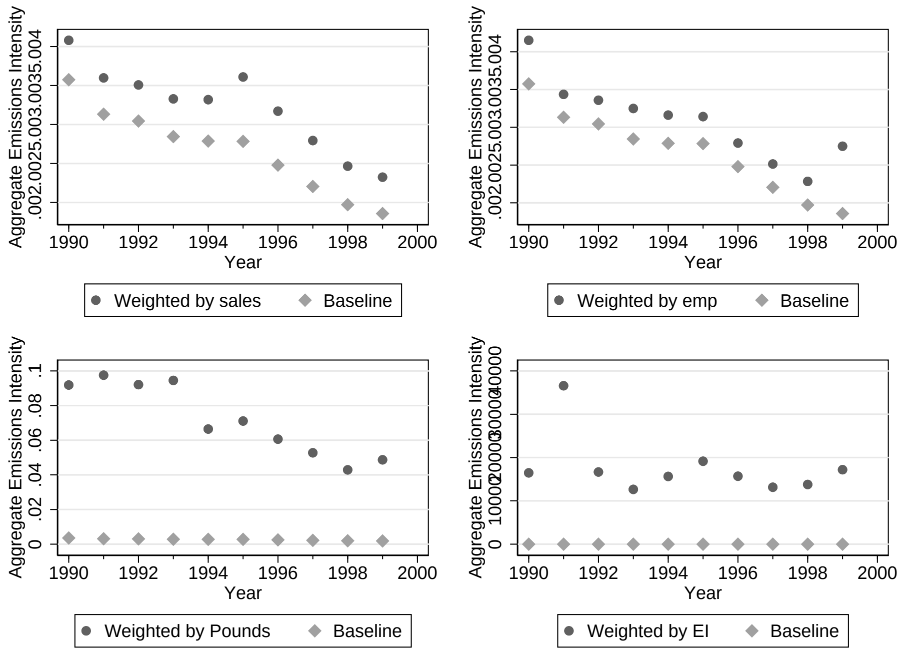

Figure 6: Change in Aggregate Emissions Intensity - Comparison of Weights

Note: The panels in this figure present the evolution of aggregate emissions intensity over time. Aggregate emissions intensity in each figure is calculated according to equation (3), the only difference being the choice of establishment weights. The "Baseline" aggregate emissions intensity included in each panel uses real establishment output weights. In the first panel, in the first row, the weight used is nominal sales from the NETS database. In the second panel, the weight is employment from the NETS database. In the first panel of the second row, the weight used is total emissions ("ModeledPounds") taken from the EPA's RSEI database, and the weight used in the final panel is emissions-intensity, calculated as the ratio of establishment emissions to real output.

completely correlate with productivity and output.[49]

Using total emissions and emissions intensity shares, in the second two panels of the figure, gives more weight to dirtier establishments. With this approach, aggregate emissions intensity appears to be decreasing at a much slower rate, suggesting that the decline in emissions may not be concentrated among the dirtiest establishments.

Alternative weighting approaches do appear to yield fruitful insights into the channels driving aggregate emissions intensity, and reveal opportunity for further research into the establishment-level factors driving observed emissions outcomes. However, due to the ability to link the results smoothly to an intuitive decomposition of aggregate emissions to better understand the establishment-level effects driving aggregate changes, we continue to prefer the output-weighting approach.

[49]This is consistent with cross-country evidence presented by Bartelsman et al. (2013) showing that there is substantial variation in the strength of the correlation between establishment productivity and size over time.

6 Conclusion

The rapid decrease in pollution emissions from U.S. manufacturers has been the subject of a great deal of attention. The existing economics literature has found that the majority of pollution emissions reductions have come from reduction in emissions per unit out output, rather than changes in the quantity or types of goods produced. In this paper, we employ a matched establishment characteristic and pollution emissions data set to confirm and extend this work. We provide evidence that the fall in criteria air pollution is mirrored by a fall in toxic pollution. As with criteria air pollutants, the primary driver of the aggregate decline in toxic pollution emissions is reduced emissions intensity within industries. Taking advantage of establishment level data, we are able to further decompose the decline in within-industry emissions intensity into four channels: reallocation among surviving establishments, entering establishments, and exiting establishments, and the within-establishment technique effect.

Our results indicate that the main driver of the cleanup in toxic pollution emissions has been within establishment improvements in emissions intensities, driving roughly 75 percent of the aggregate technique effect. However, we also demonstrate the importance of cross-establishment reallocation of economic activity towards cleaner firms and establishments that is responsible for the remaining 25 percent of the decline in aggregate emissions intensities. These improvements have markedly decreased the toxic pollution intensity of U.S. manufacturing. The differences in the contribution of the various channels across disposal media provide suggestive evidence that air pollution regulations may have played a role, but it is possible that the cleanup was driven by other environmental policies, increased competition (either domestic or foreign) or other channels either alone or in conjunction with air pollution regulations. Our result unpacks the primary channels driving the fall in toxic pollution emissions, but does not explain why there has been a remarkable improvement in within establishment environmental performance.

References

Agency, E. P. (2013). 2008 National Emissions Inventory. Technical Support Document version 3. Technical Report September.

Bartelsman, E. J., Haltiwanger, J. C., and Scarpetta, S. (2013). Cross-Country Differences in Productivity: The Role of Allocation and Selection. *American Economic Review*, 103(1):305–334.

Becker, R. A., Gray, W. B., and Marvakov, J. (2013). NBER-CES Manufacturing Industry Database: Technical Notes.

Brinkman, J., Coen-Pirani, D., and Sieg, H. (2015). Firm Dynamics in an Urban Economy. *International Economic Review*, 56(4):1135–1164.

Cherniwchan, J. (2017). *Trade Liberalization and the Environment: Evidence from NAFTA and U.S. Manufacturing*, volume 105. The Author.

Cherniwchan, J., Copeland, B. R., and Taylor, M. S. (2017). Trade and the Environment: New Methods, Measurements, and Results. *Annual Review of Economics*, 9(2):59–85.

Copeland, B. R. and Taylor, M. S. (2003). *Trade and the Environment: Theory and Evidence*. Princeton University Press, Princeton, NJ.

Cui, J., Lapan, H., and Moschini, G. (2016). Productivity, Export, and Environmental Performance: Air Pollutants in the United States. *American Journal of Agricultural Economics*, 98:447–467.

de Marchi, S. and Hamilton, J. T. (2006). Assessing the Accuracy of Self-Reported Data: an Evaluation of the Toxics Release Inventory. *Journal of Risk and Uncertainty*, 32(1):57–76.

Foster, L., Haltiwanger, J. C., and Krizan, C. J. (2001). Aggregate Productivity Growth: Lessons from Microeconomic Evidence. In Charles R. Hulten, Dean, E. R., and Harper, M. J., editors, *New Developments in Productivity Analysis*, number January, chapter 8, pages 303–372. University of Chicago Press, Chicago, IL.

Gibson, M. (2016). Regulation-induced pollution substitution. *mimeomimeo*.

Greenstone, M. (2003). Estimating Regulation-Induced Substitution: The Effect of the Clean Air Act on Water and Ground Pollution . *American Economic Review*, 93(2):442–448.

Greenstone, M., Syverson, C., and List, J. A. (2012). The Effects of Environmental Regulation on the Competitiveness of U.S. Manufacturing. *NBER Working Paper*, (18392).

Grossman, G. M. and Krueger, A. B. (1993). Environmental Impacts of a North American Free Trade Agreement. In Garber, P., editor, *U.S. Mexico Free Trade Agreement*, pages 13–56. MIT Press, 2nd edition.

Haltiwanger, J. C., Jarmin, R. S., and Miranda, J. (2013). Who Creates Jobs? Small Versus Large Versus Young. *The Review of Economics and Statistics*, 95(2):347–361.

Holladay, J. S. (2016). Exporters and the Environment. *Canadian Journal of Economics*, 49(1):147–172.

Jaffe, A. B., Peterson, S. R., Portney, P. R., and Stavins, R. N. (1995). Environmental regulation and the competitiveness of US manufacturing: what does the evidence tell us? *Journal of Economic Literature*, 33(1):132–163.

Koehler, D. A. and Spengler, J. D. (2007). The toxic release inventory: fact or fiction? A case study of the primary aluminum industry. *Journal of Environmental Management*, 85(2):296–307.

Levine, D. I., Toffel, M. W., and Johnson, M. S. (2012). Randomized Government Safety Inspections Reduce Worker Injuries with No Detectable Job Loss. *Science*, 336(April):344–347.

Levinson, A. (2009). Technology, International Trade, and Pollution from US Manufacturing. *American Economic Review*, 99(5):2177–2192.

Levinson, A. (2015). A Direct Estimate of the Technique Effect: Changes in the Pollution Intensity of US Manufacturing, 1990 - 2008. *Journal of the Association of Environmental and Resource Economists*, 2(1):43–56.

Melitz, M. J. and Polanec, S. (2015). Dynamic Olley-Pakes Decomposition with Entry and Exit. *RAND*, 42(6):1–26.

Neumark, D., Wall, B., and Zhang, J. (2011). Do Small Businesses Create More Jobs? New Evidence for the United States from the National Establishment Time Series. *Review of Economics and Statistics*, 93(1):16–29.

Olley, S. and Pakes, A. (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, 64(6):1263–1297.

Ryan, S. P. (2012). The Costs of Environmental Regulation in a Concentrated Industry. *Econometrica*, 80(3):1019–1061.

Shapiro, J. S. and Walker, R. (2016). Why is Pollution from U.S. Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade. *mimeo*.

Shapiro, J. S. and Walker, R. (2017). Why is Pollution from U.S. Manufacturing Declining? The Roles of Trade, Regulation, Productivity, and Preferences. *mimeo*.

Appendix

7 Data Set Construction

In this section we provide a brief description of the process we used to construct the matched polluting and establishment characteristics datatset. The initial dataset was constructed for (and described in) Holladay (2016). We begin by matching as many RSEI establishments to the NETS database as possible. The initial match was based on DUNS number, a proprietary establishment-specific identifier created by Dunn and Bradstreet. Unfortunately, the EPA treats DUNS number as optional in TRI submissions, so many polluter records do not have a DUNS. For these establishments we used a a fuzzy matching procedure based on common fields in the NETS and TRI: establishment name, address, latitude-longitude and industry. We also acquired a 10% sample of all manufacturing establishments in the NETS which we use to construct comparisons to other measures of manufacturing output.

We merge the combined TRI-NETS data with the NBER-CES manufacturing database. That database contains industry-year specific price indices. We use those price indices to deflate the output reported in the NETS. We create emissions intensity measures by dividing an establishment's reported TRI emissions (and hazard) by the deflated output. The NETS allows multi-product establishments to report multiple SIC codes. Where reported industries differ between the datasets we defer to the RSEI.

While the TRI and NETS include data on non-manufacturing industries, we restrict our attention to establishments that report SIC 20-39 as their primary classification. We drop from the analysis six sectors (logging and publishing sub-industries: SICs 2411, 2711, 2721, 2731, 2741 and 2771) that were fully reclassified out of manufacturing over the sample period, because they are missing price-index data across the entire sample. We follow Levinson (2009) in dropping observations from 9 computer-related sectors whose price indices declined over the sample period: electronic computers (3571), computer storage devices (3572), computer terminals (3575), computer equipment n.e.c. (3577), calculating and accounting machines (3578), household audio and video equipment (3651), telephone and telegraph apparatus (3661), semiconductors (3674), and magnetic and optical media (3695) because they experience deflation across our sample period. We focus on the years 1990-2000 because the output data reported by the NETS (both nominal and real) are not consistent with those reported in the NBER-CES after 2000. The NETS data contains a combination of establishment reported output, Dunn and Bradstreet imputed output and NETS imputed output. Holladay (2016) describes in detail how the output data is imputed and provides simple robustness checks for the imputed data.

We keep all observations for which emissions intensity data is not null, which in effect drops observations without emissions data and drops observations without real output data.[50] The final data set consists of 92,210 establishment-year observations. We create firm entry and exit variables using the NETS "firstyear" and "lastyear" variables across the "HQdunsnumber" variable. Thus, we capture any entry or exit of establishments within a given firm that is not entering or exiting in the changes due to surviving establishments. It is important to note that our definition considers reallocation among surviving establishments within surviving firms to be "reallocation"–this definition only affects how we define entry and exit–which we define as a firm-level phenomena.

Finally, we decompose annual emissions changes according to process described in paper theory section.

8 Matching Establishment Characteristics and Pollution Emissions

In this section we briefly compare our matched TRI-NETS data to the full NETS and TRI samples. Table 4 compares entry and exit rates of firms in the matched TRI-NETS sample to the Census of Manufactures. We observe lower rates of entry and exit in the NETS than is reported in the Census data. Haltiwanger et al. (2013) makes a similar observation for a larger extract from the NETS, and suggests that this may be due to the NETS database having less coverage of very small establishments, particularly those with one or zero employees.

However, since our focus is the role of entry and exit on emissions outcomes, the fact that entry and exit in our sample may not exactly match the broader economy is of less concern for two reasons. First, by comparing the sales of polluters to sales of all establishments in our sample, we observe that establishments that report emissions tend to be larger.[51] Thus, the matched data over-samples large establishments relative to the broader economy, but this merely reflects the fact that these establishments emit the vast majority of toxic pollution. Second, the distribution of emissions in our matched TRI-NETS sample is consistent with the pollution measures in the complete TRI data set, lending further confidence that the results we report are representative of the channels driving aggregate emissions for all TRI reporters.

[50] In the next section we compare the composition of the final data set to the raw NETS, TRI and NBER-CES datasets to evaluate the effect of these dataset construction choices.

[51] Relatively large establishments tend to enter and exit at lower rates than their smaller counterparts, but without access to the establishment level Census of Manufacturers data we cannot confirm whether the entry and exit rates in the NETS are consistent with broader measures for large manufacturing establishments.

Table 4: Market Share, Entry and Exit Rates by Firms (NETS)

	Market Share		Entry and Exit Rates			
	NETS		NETS		US Census	
Year	Entrants	Exiters	Entry	Exit	Entry	Exit
1991	0.2	14.3	0.7	10.8	7.8	7.8
1992	0.3	5.1	0.8	5.6	7.1	8.1
1993	0.4	4.5	2.3	5.5	8.3	7.6
1994	0.6	1.4	1.3	3.1	8	7.6
1995	0.2	5	1.4	6.7	8.2	7.4
1996	0.3	3.9	1.5	7	7.8	7.5
1997	0.4	2.5	1.6	4.4	7.4	7.5
1998	0.4	3.6	1.2	5.2	7.1	7.5
1999	0.4	1.7	1.4	2.6	6.2	7.6
Mean	0.4	4.7	1.4	5.7	7.5	7.6

Entrants and Exiters are identified by a firstyear and lastyear indicator (respectively) in the NETS database and their HQDuns number. Market Share is the ratio of total output by establishments in each group (Entrants or Exiters) to total output in each year. Entry Rates are the ratio of the number of entering firms to the total number of firms. Exit Rates are the ratio of the number of exiting firms to the total number of firms in the previous year.

As we note in the main text, establishments reporting emissions tend, on average, to be slightly larger than the average establishment in the larger NETS sample. Table 5 reports the sample statistics for sales (and log sales) reported in the full NETS sample as compared to the the sub-sample of polluters. Just over five percent of NETS establishment observations are matched to TRI reporters. The larger NETS sample contains a number of establishments that do not report emissions because they do not pollute or do not meet the reporting requirements. Because establishments with fewer than 10 employees are exempt from reporting to the TRI we expect to see fewer small establishments in the matched sample and that is consistent with the statistics reported in Table 5.

Figure 7 reports the distribution of log sales across the two data sets. The matched data set is in the top panel and the full NETS data is in the bottom panel. The shape of the distribution is quite similar, but, consistent with the summary statistics reported above, the matched sample is shifted roughly three log points to the right.

As reported in Holladay (2016), the matched data set contains around two-thirds of the full set of TRI reporting manufacturers. Figure 8 reports the distribution of log pounds of toxic emissions for the matched TRI-NETS set (top panel) and the EPA's full TRI data set

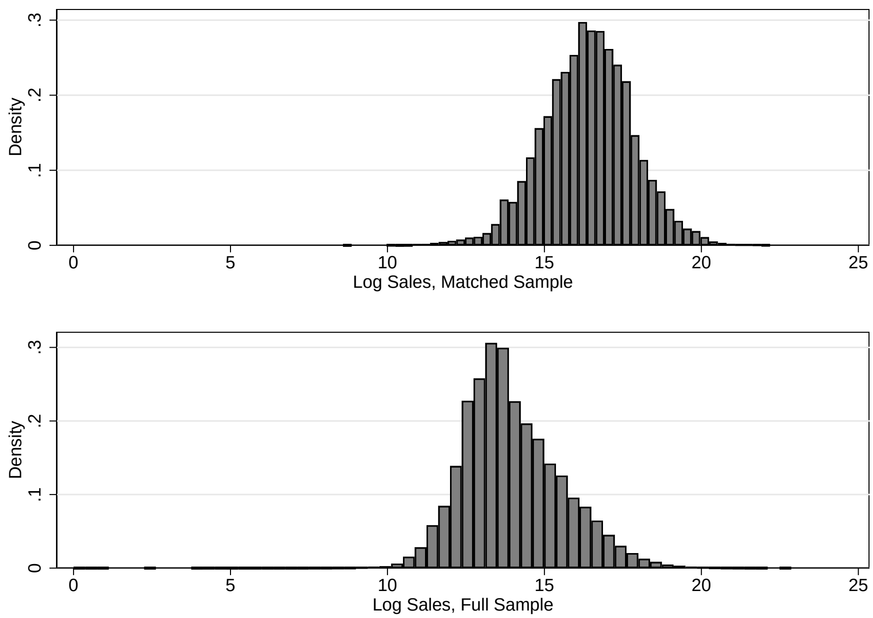

Figure 7: Sample Comparison: NETS (Sales)

Note: The top panel shows the distribution of establishment log-sales of polluting establishments from the NETS database that also report emissions to the TRI. The lower panel shows the distribution of log-sales for the larger 10-percent NETS sample, including establishments that do not report emissions to the TRI. Both distributions are restricted to manufacturing (SIC2 20-39) establishments operating from 1990 - 2000. The establishment sales data has been weighted by sector-specific NBER-CES price indices.

Table 5: Sample Comparison: NETS (Sales)

	Sales		Log Sales	
	Polluters	Both	Polluters	Both
Mean	36,900,000	7,565,452	16.4	14.1
Std. Dev.	104,000,000	47,100,000	1.4	1.6
Median	13,500,000	1,050,000	16.4	13.9
Min	5,336	0	8.6	0
Max	4,330,000,000	8,490,000,000	22.2	22.9
Obs.	92,210	1,314,619	92,210	1,314,618

The first two columns compare the sample Sales statistics for establishments matched from the RSEI to the NETS database and the larger 10-percent NETS sample including both polluters and non-polluters. The second two columns compare the sample Log-Sales statistics.

(bottom panel). The distributions are very similar. The sample mean, standard deviation and median are within rounding error of each other. We are confident that the matched sample is a good representation of the full TRI data set.

The set of establishments that report to TRI does not represent the full distribution establishments in the NETS. However, because our focus is on the environmental performance, we are satisfied that the merged data provides an accurate representation of the channels driving the evolution of toxic pollution emissions that are reported by TRI.

9 Defining entry and exit

As noted in the main text, the definition of entry and exit has important implications for the estimation of the impacts of entry, exit and reallocation. In this section we briefly present decomposition results when we define entry and exit as occurring at the establishment level, without regard to the status of the parent firm. These results indicate that manufacturing establishments are reallocating production towards dirtier establishments, and that a tremendous amount of the observed cleanup in toxic pollution emissions is due to the within establishment cleanup, the technique effect. The finding of a positive reallocation effect is reversed when we define entry and exit as occurring at the firm level. Table 6 reports the reallocation and within establishment technique effects using the establishment level definitions of entry and exit.

Reallocation of resources within industries tends to be towards dirtier establishments, driven primarily by the exit of relatively clean establishments. This is offset by the relentless cleanup in surviving establishments which drives the overall cleanup in the manufacturing

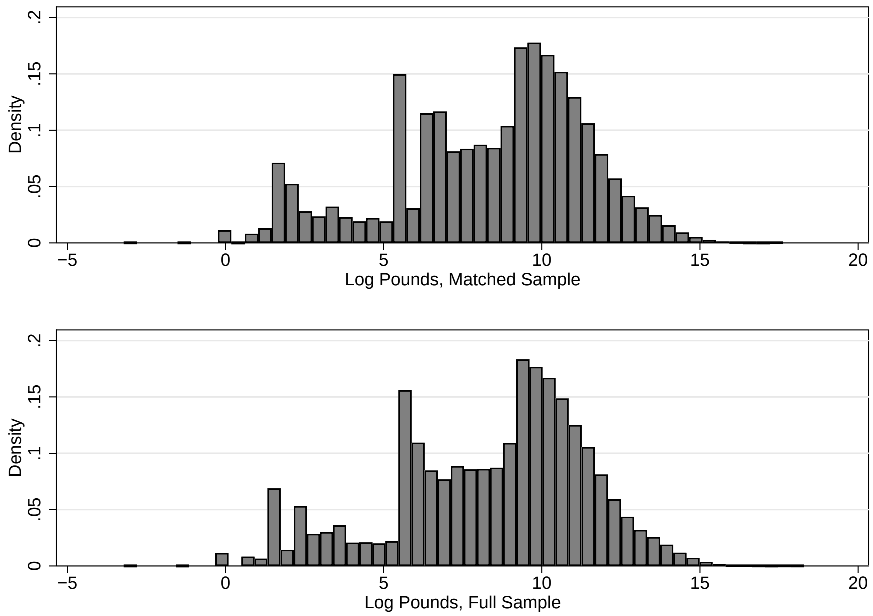

Figure 8: Sample Comparison: RSEI (Pounds)

Note: The log-emissions data for the top panel of this figure are taken from matched RSEI-NETS establishment-level data set described in the text. The distribution in the lower panel is based on the log-pounds of manufacturing establishments (SIC2 20-39) from the full RSEI sample.

sector. The results indicate that the cleanup is not being driven by environmental regulation forcing old, pollution intensive manufacturing establishments out of business.

Using this broader definition of entry and exit, exiting establishments tend to be cleaner than the incumbents they leave behind. As a result, this reallocation channel raises toxic emissions from the manufacturing sector during our sample period. The decline in the aggregate technique effect masks the fact that exit of relatively clean establishment raises emissions by between 1 and 3.7 percent depending on the year. To better understand this result we aggregated the effect of exiting establishments on total manufacturing emissions by industry and year. Industrial organic chemicals (SIC 2869) experiences the most exit from relatively clean establishments, but paper mills (SIC 2611) and pulp mills (SIC 2621) are not far behind. Only a handful of manufacturing industries (29 out of 450 in our sample) get cleaner due to exit of relatively dirty establishments. Only for wet corn milling (SIC 2046) and household cooking equipment (SIC 3631) does exit of relatively dirty establishments contribute more than 0.1 percent to the overall clean up.

By using these two definitions together the findings give some additional insight into the operations of manufacturing firms and their environmental consequences. In particular, our results indicate that a substantial portion of the aggregate technique effect is due to firms merging cleaner operations with dirtier, and not merely from the installation or adjustment

Table 6: Change in Aggregate Emissions Intensity by Channel

	(1)	(2)	(3)	(4)	(5)	(6)
	Surviving Establishments			Entering Establishments	Exiting Establishments	Total (All Establishments)
Year	Total	Within	Between			
1991	-14.7	-21.6	6.9	-0.9	3.0	-12.6
1992	-4.6	-16.6	12.0	-1.4	3.2	-2.9
1993	-12.3	-12.2	-0.2	-2.0	3.6	-10.8
1994	-3.5	-3.3	-0.2	-0.3	3.4	-0.5
1995	-4.4	-7.7	3.4	-1.8	2.3	-3.9
1996	-13.6	-14.4	0.8	-1.0	3.3	-11.3
1997	-11.5	-7.4	-4.0	-1.2	3.9	-8.8
1998	-12.4	-12.4	0.0	-0.9	2.8	-10.6
1999	-9.6	-3.2	-6.4	-0.7	2.7	-7.6
2000	-5.6	-3.5	-2.1	3.7	2.9	1.0

Note: Change in toxic pollution emissions attributable to surviving, entering and exiting establishments. Change in emissions in surviving establishments are further decomposed into changes within establishments in column 2, which we term the "pure technique" effect and reallocation among survivors in column 3. Units are percent of 1990 toxic release inventory emissions in our sample.

of production technologies at the dirtier establishments.