Independent Statistics & Analysis
U.S. Energy Information Administration

The Energy Efficiency Gap and Energy Price Responsiveness in Food Processing

October 2020

This report was prepared by the U.S. Energy Information Administration (EIA), the statistical and analytical agency within the U.S. Department of Energy. By law, EIA's data, analyses, and forecasts are independent of approval by any other officer or employee of the U.S. Government. The views in this report should not be construed as representing those of the U.S. Department of Energy or other federal agencies.

The Energy Efficiency Gap and Energy Price Responsiveness in Food Processing

The National Energy Modeling System (NEMS) is an integrated energy forecasting tool used by the U.S. Energy Information Administration (EIA) to prepare the *Annual Energy Outlook* and other analytic products. NEMS includes modules for the energy supply sectors as well as sources of energy demand. In particular, the NEMS Industrial Demand Module (IDM) estimates energy consumption by energy source (including fossil fuel used as a feedstock) for 15 manufacturing and 6 non-manufacturing industries.

One of these manufacturing industries is food processing. Although food processing is less energy intensive than sectors such as steel or cement, it is still considered energy intensive in the NEMS IDM because of its relatively large energy consumption (1.1 quadrillion British thermal units of energy according to EIA's 2014 *Manufacturing Energy Consumption Survey*).

The IDM uses an end-use approach for modeling the food processing industry. End-use models use a capital stock vintage accounting framework that models energy use in new additions to the stock and in the existing stock. The model assumes that energy intensity decreases for both but that energy intensity of existing stock declines more slowly. Specifically, new additions employ state-of-the-art technologies to bring down energy intensity at a given rate, while existing stock undergoes retrofitting and replacement of equipment from normal wear and tear, which incorporates 50% of the improvement that is achieved by adding new stock.

By design, the existing end-use approach uses the observed average efficiency across firms in its estimation process. To understand potential future changes in energy intensity, one approach seeks to assess the *energy efficiency gap*—generally defined as the difference between the actual level of efficiency and an estimated, practically achievable level.

To better understand this, EIA commissioned Leidos, Inc., to prepare the following report *The Energy Efficiency Gap and Energy Price Responsiveness in Food Processing*. The report provides an analysis of this efficiency gap, as well as elasticities calculated in the stochastic frontier energy demand estimation process, which allow EIA to estimate how much of an efficiency improvement may be available in the food processing industry. In simulating efficiency improvements (by setting the target level of efficiency based on a percentile of the cumulative efficiency distribution), the estimate of potential reduction in energy use is by definition smaller but empirically more reasonable. The report finds that a practical improvement modeled by bringing the bottom half of the distribution up to the median can result in a 13.0% reduction in the efficiency gap when averaged across all models. If the lower end of the distribution is brought up to the ENERGY STAR® rating of the 75th percentile, the reduction in the efficiency gap increases to an average of 15.5%, and when targeting the 95th percentile, the reduction in the efficiency gap increases to 20.0%.

When referencing the report, cite it as a report by Leidos, Inc., prepared for the U.S. Energy Information Administration.

The Energy Efficiency Gap and Energy Price Responsiveness in Food Processing[1]

Gale Boyd[2] and Matt Doolin[3]

Duke University

May 19, 2020

Abstract

This paper estimates stochastic frontier energy demand functions with non-public, plant-level data from the U.S. Census Bureau to measure the energy efficiency gap and energy price elasticities in the food processing industry. The estimates are for electricity and fuel use in 4 food processing sectors, based on the disaggregation of this industry used by the National Energy Modeling System Industrial Demand Module. The estimated demand functions control for plant inputs and output, energy prices, and other observables including 6-digit NAICS industry designations. Own price elasticities range from -0.6 to -0.9 with little evidence of fuel/electricity substitution. The magnitude of the efficiency estimates is sensitive to the assumptions but consistently reveal that few plants achieve 100% efficiency. Defining a "*practical level of energy efficiency*" as the 95th percentile of the efficiency distributions and averaging across all the models result in a ~20% efficiency gap. However, most of the potential reductions in energy use from closing this efficiency gap are from plants that are "low hanging fruit"; 13% of the 20% potential reduction in the efficiency gap can be obtained by bringing the lower half of the efficiency distribution up to just the median level of observed performance. New plants do exhibit higher energy efficiency than existing plants which is statistically significant, but the difference is small for most of the industry; ranging from a low of 0.4% to a high of 5.7%.

[1] This paper was funded by the U.S. Energy Information Administration (EIA), Office of Energy Consumption and Efficiency Analysis while working with Leidos Inc. under contract with EIA (DE-EI0000564) and was prepared while the authors were Special Sworn Status researchers at the Triangle Research Data Center, a member of the Federal Statistical Research Data Center Network under FSRDC Project Number 2173. An earlier draft benefited from comments by Kelly Perl and Peter Gross. Any errors or omissions are the responsibility of the authors. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau or the EIA. All results have been reviewed to ensure that no confidential information is disclosed.

[2] Corresponding author: Associate Research Professor, Social Science Research Institute, Duke University, Box 90989, Durham, NC 27708. E-mail: gale.boyd@duke.edu.

[3] Research Associate, Social Science Research Institute, Duke University, Box 90989, Durham, NC 27708. E-mail: matthew.doolin@duke.edu

Introduction

The notion that energy demand suffers from the “energy paradox” or the “energy-efficiency gap” is pervasive in the energy demand literature. (Gerarden, Newell et al. 2015) define the “energy paradox” as ‘*the apparent reality that some energy-efficiency technologies that would pay off for adopters are nevertheless not adopted. This basic definition relates to the issue of private optimality.*’ This paper uses the term “energy efficiency gap,” in the narrower private notion that (Gerarden, Newell et al. 2015) term the energy paradox.

An exhaustive review of the literature on the energy efficiency gap is beyond the scope of this paper; some examples include (Howarth and Andersson 1993, Jaffe and Stavins 1994, Huntington 1995, Allcott and Greenstone 2012, Boyd and Zhang 2013, Boyd and Curtis 2014, Boyd 2016). This literature can loosely be separated into two threads: identification and quantification. The first thread of this literature, reviewed by (Gerarden, Newell et al. 2015), focuses on identifying potential reasons for the gap, e.g. lack of information, market failures, etc. Much of the quantification literature employs an engineering economics perspective. (Worrell, Ramesohl et al. 2004) review modeling industrial energy demand from this engineering economics perspective.

The conservation supply curve (CSC) is popular representation of the engineering approach to quantification and is typified by (Rosenfeld, Atkinson et al. 1993). Engineering economics is not the only approach to the CSC. (Huntington 1995) discusses the intellectual connections between the energy efficiency gap literature and the literature on measuring production efficiency. (Blumstein and Stoft 1995, Stoft 1995) make the connections between production economics and the CSC more explicit by showing how the CSC can be constructed from the production function when inefficiency is present. This paper presents a quantification of energy efficiency, but not with an engineering economics approach. It draws on economic production theory and statistical modeling as the basis for quantifying the energy efficiency gap by applying stochastic frontier analysis (SFA). It is therefore more in the spirit of (Huntington 1995) with (Boyd 2008, Filippini and Hunt 2016) being empirical examples.

The National Energy Modeling System (NEMS) model is an integrated energy forecasting tool used by the U.S. Energy Information Administration to prepare the Annual Energy Outlook. NEMS includes modules for the energy supply sectors as well as sources of energy demand. In particular, the NEMS Industrial Demand Module (IDM) estimates energy consumption by energy source (including fossil fuel used as a feedstock) for 15 manufacturing and 6 non-manufacturing industries. One of the manufacturing sectors is food processing. While food processing is less energy intensive than sectors like steel, cement, etc., it is still considered energy intensive for the purposes of the NEMS IDM. Unlike other energy intensive sectors which employ a process flow modeling approach, the food processing industry is modeled using an end use approach. The end use approach does not currently account for the possibility of an energy efficiency gap. The difficulty in representing the efficiency gap is not an inherent weakness of the end use approach, but requires additional information. This analysis provides the additional information that can be incorporated into the IDM to model the energy efficiency gap in food processing.

Background on energy use in food processing

The Food Processing Industry, as defined by the North American Industry Classification System (NAICS) code 311,

"… transform livestock and agricultural products into products for intermediate or final consumption. The industry groups are distinguished by the raw materials (generally of animal or vegetable origin) processed into food products. The food products manufactured in these establishments are typically sold to wholesalers or retailers for distribution to consumers, but establishments primarily engaged in retailing bakery and candy products made on the premises not for immediate consumption are included."[4]

The 2014 Manufacturing Energy Consumption Survey (MECS) reports that 1.1 quadrillion British Thermal Units (Btu) of energy was consumed in NAICS 311[5].

Approximately 66% of fuel use in NAICS 311 is natural gas for heat and power while 30% of fuel use is coal and bio-waste. The use of the latter two fuel types is concentrated within two NAICS codes: 311221 (Wet corn milling) and 31131 (Sugar refining). IDM uses four groupings for food processing. This study focuses on those sectors, but with a key modification. NAICS 311221 and 31131 were disaggregated from the parent industry IDM grouping such that the analysis in this paper is conducted using 6 industry groupings, rather than the 4 used by the IDM. The IDM groups and the groups used in this study are shown in Table 1.

Table 1 Industry grouping used in the study, defined via NAICS code

IDM industry groups	
IDM-1	3112
IDM-2	3115
IDM-3	3116
IDM-4	Balance of 311
Study industry groups	
IG-1	IDM-1 less 311221 (wet corn milling)
IG-2	IDM-2
IG-3	IDM-3
IG-4	IDM-4 less 31131 (sugar processing)
IG-5	311221 (wet corn milling)
IG-6	31131 (sugar processing)

Overview of the approach

A stochastic frontier regression analysis (SFA) is applied to pooled cross sections using plant level data from the quinquennial Census of Manufacturing (CMF) for the years 1992, 1997, 2002, 2007, and 2012. Using the SFA estimates, the study measures the efficiency of the plant by computing the ratio of the predicted frontier level of energy use divided by the actual energy use. For example, if that ratio is 80% then energy use could be reduced by 20%. This ratio is our measure of "*efficiency*" and the "*efficiency gap*" is one minus the ratio. The efficiency of plants entering the industry can be compared to those who are already in operation in that same year. This ratio measure of efficiency, based on the estimated frontier, is then modified to take into account the empirical distribution of efficiency to provide a

[4] https://classcodes.com/lookup/naics-3-digit-subsector-code-311/

[5] All energy data are from the 2014 *MECS Table 1.2 First Use of Energy for All Purposes (Fuel and Nonfuel*). These data are measured at end-use; i.e. electric generation losses are not included.

"*practical measure of efficiency*". This modification is explained further in the results section. The methodology and data construction used here is comparable to two other plant level studies (Boyd and Lee 2019, Boyd and Lee 2020).

To assess the sensitivity of the efficiency estimates to model specification, two approaches are presented. The first is a common frontier that is estimated using pooled data across all of the subsectors, but allowing for the possibility that the distribution of energy efficiency varies across those sectors using various distributional assumptions. The second approach is estimation of individual frontiers for each sector using a two – stage estimation method similar to that used by (Boyd and Lee 2020) who incorporated instrumental variables to account for electricity price endogeneity in the chemical industry. Initial analysis of this industry did not find any evidence of electricity price endogeneity based on a Hausman test, so the instrumental variable approach was not used[6]. (Boyd and Lee 2019) found the same lack of price endogeneity in the study of metal-based durables (MBD), so food processing seems similar to MBD as opposed to chemicals, i.e., electricity prices payed by firms in these food processing and MBD exhibit little endogeneity. It is worth noting that chemicals is very energy intensive, food processing less intensive, and MBD rather low.

The report is organized as follows. The first section describes the data sets used and transformation required to get the final variables needed for the analysis. The second section describes the general specification of the model, the pooled frontier approach and the two-stage, industry specific approach to estimate the frontier(s) for persistent and time-varying efficiency. The third section presents the frontier parameter estimates, the efficiency of new firms, and a simulation of efficiency gains based on achieving different levels of efficiency using the percentiles of the estimated distributions.

Data

Data for the study are non-public plant-level Census Bureau data made available at the Triangle Federal Statistical Research Data Center. These data are protected under Title 13 and 26 of the US Code and used with permission from the Bureau. The data sources are the MECS and the Census of Manufacturing (CMF). MECS is a sample based survey conducted in 1985, 1988, 1991, 1994, 1998, 2002, 2006, 2010, and 2014[7]. The CMF is part of the quinquennial Economic Census; it includes all establishments operating during the analysis time period, 1992, 1997, 2002, 2007, and 2012. Both data sources span similar time periods but, for the most part, different years. The MECS and CMF each have advantages and disadvantages.

Data needed for the analysis include energy use and prices along with production activities and other location specific attributes. While the MECS provides the most detailed data on energy use, particularly cost and quantity of fuels by type, the MECS is a stratified sample and not a balanced panel so the presence (absence) of an observation is not an indicator of entry (exit) in the industry. We need this information on entry/continuing status to measure the relative efficiency of entering vs continuing plants. Using the CMF solves this problem.

---

[6] Plant level electricity prices are available in the CMF data, but not fuel prices, so the examination of potential prices endogeneity was limited to electricity prices only.
[7] 2014 was the most recent year, but not yet available to external researchers.

The availability of plant level electricity use and prices in the CMF is one advantage of this data set. The CMF provides plant level electricity consumption and costs, from which a plant level average price can be computed directly. (Davis, Grim et al. 2012) analyze the dispersion of those prices in detail. However, the CMF only reports expenditures on fuels, not quantities, so Btu fuel consumption is imputed from fuel costs in the CMF assuming state level average price of fossil fuels. (Boyd 2016, Boyd and Lee 2019, Boyd and Lee 2020) impute fossil fuel use from the state average industrial price of natural gas. This was seen as a reasonable assumption for the metal based durables industries, because publicly available MECS data from 2010 for these 5 sectors suggests that 88% to 98% of the purchased fuel in this sector is natural gas. This is less true for food, with 62% of fuel use reported in MECS is natural gas. However, in food processing, 30% of fuel use is coal and bio-waste. This is concentrated in two NAICS codes; 311221 – *Wet Corn milling* and 31131 – *Sugar Refining*. These are sub-sectors of 3112 *Grain and Oilseed Milling* and Other – *Balance of 311*, respectively. For this and other reasons described below, NAICS 311221 and 31131 were therefore disaggregated from the parent industry grouping such that much of the analysis in this paper was conducted using 6 industry sector groupings rather than the four groups used by the IDM.

There is no simple "best choice" regarding using CMF vs MECS data for this study, particularly with respect to fuel use, but also industry coverage. In terms of coverage, the CMF is a census and covers virtually all plants in the industry[8], but the CMF does not have fuel consumption, only expenditures. The initial approach taken by this study was to impute Btu fuel consumption by taking fuel expenditures reported in the CMF and dividing by the state level industrial natural gas prices as published by the EIA's State Energy Data System (SEDS). This approach is not without its problems, but has been used in other studies. The well-known caveat is that the reported industrial natural gas price from the SEDS reflects only the portion of natural gas sold to industrial users who get their gas through a local distribution company (LDC). Smaller food processing plants likely get their gas from an LDC, but very large plants may have a direct pipeline arrangement.

The MECS provides detailed data on energy use, specifically cost and quantity of fuels by type, so does not require imputation. The MECS is a stratified sample that samples large plants in energy intensive sectors with certainty. Some energy intensive industries in certain 6-digit NAICS codes have higher sampling rates that allow for those industries to be published separately. Only 311221 – *Wet Corn milling* and 31131 – *Sugar Refining* fall into this category for food processing, although some 4-digit NAICS are sampled at rates that allow aggregate publication. What this implies for this study's use of the micro data it that only a few industries will have sufficient observations for our analysis. The solution is to combine the CMF data with the imputed fuel use with select data from plants in the MECS sample, focusing on those sectors with higher rates of sampling. While MECS is a sample, its primary advantage is in the fossil fuel detail in sectors that have higher sampling rates. This study leverages this to obtain a better estimates of the fuel quantities in NAICS 311221 and 31131, where fuel use is more likely to be something other than natural gas.

One major down side of the imputation approach used with the CMF fuel expenditure data and SEDS prices is that natural gas is more expensive than coal and other, i.e. waste derived, energy forms. If a plant is using coal then the imputation approach will bias the fuel estimate downward, making plants look more efficient. For NAICS 311221 and 31131 we leverage the detail in MECS to improve on the

[8] Very small plants with less than 5 employees may have their data imputed from administrative records. Those observations are not used here.

imputation approach. First we identify plants in the CMF that are also in MECS. Since 2002 is the only year when CMF and MECS coincide, we pick the closest MECS year for all other CMF years. For example, the 2007 CMF is paired with the 2008 MECS, the 2012 CMF with the 2014 MECS, etc. Taking 2007 as an example, since we need an energy estimate for 2007 we take the 2008 MECS fuel quantity (in MMBtu) and scale it by the ratio of total value of shipments in the CMF and MECS, adjusted for inflation[9]. This simplifying assumption that energy use is proportional to sales isn't perfect, but is more likely to avoid the downward bias using cost data and natural gas prices. Published MECS data confirms that many plants in NAICS 311221 and 31131 use coal and other fuels. This approach prevents the efficiency estimates from being overstated in these sectors.

One additional piece of information that ties electricity and fuel is the practice of combined-heat-and – power (CHP). Plants that generate some of their own electricity, not uncommon in this industry, will likely purchase more fossil fuel and less electricity. To account for CHP, the ratio of on-site generated power to total net electricity consumption, defined as purchased electricity + self-generated electricity – electricity sales to the grid, is computed for all plants that report CHP. This ratio is used in the demand model to control for the decision to use CHP, without which the frontier demand equation approach would generate biased estimates. It is worth pointing out that the decision to employ CHP is driven, in part, by relative electricity and fuel prices. The cross price elasticities in the demand models will not capture the CHP decision, since we control for CHP as an exogenous factor. While the data used here do provide insights into CHP in this and other industries, an integrated study of energy efficiency and CHP would require additional analysis.

Plant level shipment values, adjusted for inventory changes, are used to measure production. Labor is measured in production worker hours. Capital stock is computed in the micro data using perpetual inventory methods on investment data, separately for both plant and equipment and then aggregated (Foster, Grim et al. 2016). Non-energy materials are computed by subtracting total material expenditures less expenditures for electricity and fuels. All data in dollar values are deflated using implicit GDP price deflator. Since weather may impact energy use, data on annual heating and cooling degree days based on the first three digits of the plant zip code is used based on data in the Energy Star Portfolio Manager Degree Day Calculator[10].

Methodology

This section briefly presents the ad-hoc demand model specification. The demand model is "ad-hoc," since it is not a separable system of demand equations derived from an underlying cost function structure, as described by (Caves, Christensen et al. 1982). (Bardazzi, Oropallo et al. 2015) is one recent example of such a cost function based system of demand equations. This ad-hoc demand approach is implemented by adding energy prices to the energy factor requirement function described by (Boyd 2008), which is equivalent to a directional input distance function. (Boyd and Lee 2019) motivate this by considering the energy prices as a modification of the direction of the distance function, but do not model that connection explicitly. A review of stochastic frontier applications for energy use can be found in (Filippini and Hunt 2015).

[9] The CMF has the additional production details on labor, materials, and capital stock. That is why we use MECS to estimate fuel use for the appropriate CMF data years.

[10] https://portfoliomanager.energystar.gov/pm/degreeDaysCalculator

This study compares two different approaches to estimating the frontier demand model. The first is a common frontier equation, allowing the distributional assumptions to vary by industry grouping. The second treats each industry grouping separately and applies a two stage frontier approach proposed by (Kumbhakar, Lien et al. 2014) which allows for the decomposition of efficiency into two components; one is plant specific and constant over time (persistent efficiency) and one that is time varying (transient efficiency). The details of these two approaches are described below.

Common frontier with sector specific distributions

Following (Boyd and Lee 2019) we specify the energy demand equation for the two primary energy types in each of the six industry sectors. We consider log linear models of the form,

$$lnE_{j,i,t} = f(\,lnY_{i,t}, lnK_{i,t}, lnL_{i,t}, lnM_{i,t}, lnPrE_{i,t}, lnPrF_{i,t,s},$$

$$CDD_{i,t}, HDD_{i,t}, DNAICS \;\;, \; DYear_t, GERATIO_{i,t}) \; + \; v_{j,i,t} \; + \; u_{j,i,t} \qquad (1)$$

Where

$lnE_{j,i,t}$	=	ln of energy use
$lnY_{i,t}$	=	ln of production or output
$lnL_{i,t}$	=	ln employment or other measure of labor
$lnK_{i,t}$	=	ln capital stock
$lnM_{i,t}$	=	ln of non-energy material use
$lnPrj_{i,t,s}$	=	ln price of energy j = E, F [11]
$CDD_{i,t}$, $HDD_{i,t}$	=	cooling or heating degree days
$DNAICS$	=	dummy for the 6-digit NAICS code for plant *i* (not reported in the results)
$DYear_t$	=	dummy for the survey year (not reported in the results)
$GERATIO_{i,t}$	=	ratio of self-generated electricity to total purchased + generated - sold

j = energy type (purchased electricity and fossil fuel, onsite renewables, bio-waste, etc. are excluded)
i = individual establishment (i.e., manufacturing plant)
s = state
t = year of the observations i.e., 1992, 1997, 2002, 2007, and 2012

The standard SFA approach is to treat the error term as the sum of two terms representing statistical noise, $v_{j,i,t}$, and inefficiency, $u_{j,i,t}$, respectively. The distribution of u_i may be assumed to be half normal, truncated normal, or exponential.[12]

$$u_i \sim N^+(0, \sigma_u^2) \quad (2a)$$

$$u_i \sim N^+(\mu, \sigma_u^2) \quad (2b)$$

[11] The subscript 's' refers to state level, but we use both state and plant level prices as detailed below.

[12] In the SFA literature the truncation point is assumed to be at zero since efficiency is one-sided. This is denoted N+. Half normal is assumed to be normal with a 0 mean and truncated at zero, but only the upper half, truncated normal is still truncated at zero (N+) but a non-zero mean. This changes the shape of the distribution in a much more flexible way. In practice truncated normal is harder to estimate since the distribution has two parameters instead of one.

$$u_i \sim Exp(\theta) \qquad (2c)$$

To illustrate how these distributional assumptions can vary to reflect the underlying empirical distribution of efficiency, some examples are shown in Figure 1. The half normal represents most plant as being near the frontier, but the efficiency distribution declines more rapidly than the exponential. The truncated normal, depending on the location of the mean, may have a mode that is similar to the half normal and exponential (if the mean is negative), but can also capture if relatively more plants are inefficient, i.e. have a non-zero mode that is not near the frontier, shown in the example with μ=1, σ=0.5. The log form of the models imply that the efficiency distributions are percentage deviations from the estimated frontier, not absolute deviations. The log form of the model also means that all parameter estimates, except for CDD and HDD which are not in logs[13], are directly interpreted as the corresponding elasticity.

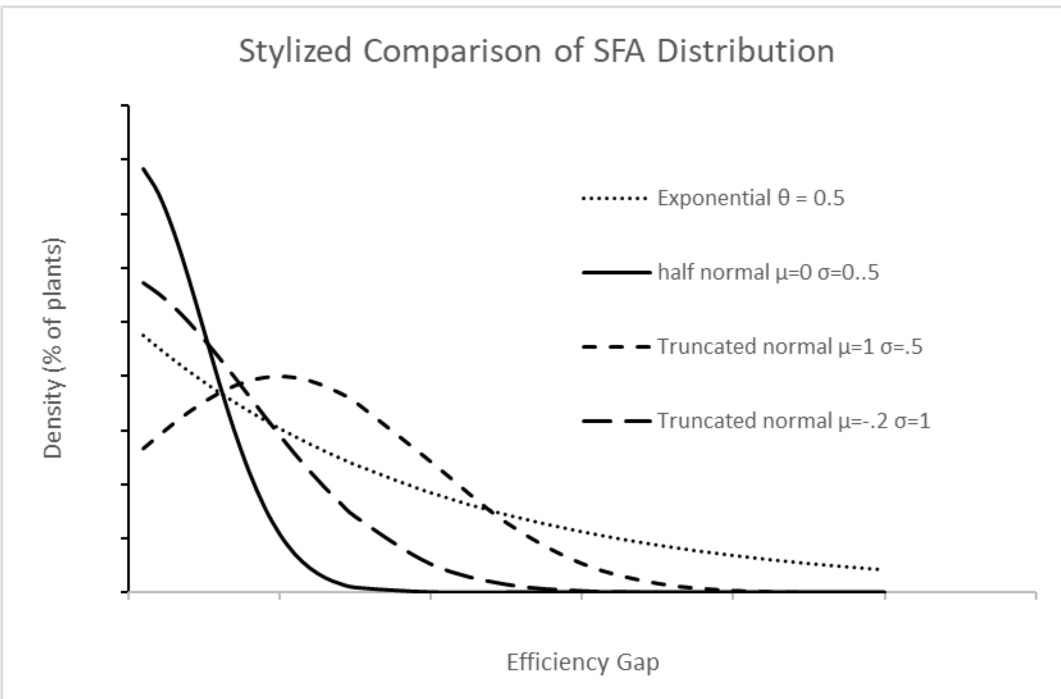

Figure 1 Comparison of SFA assumptions

This framework can be extended by assuming that the parameters of the efficiency distributions exhibit systemic differences, in this case that they vary by industry grouping.

$$\sigma^2_{u,i} = \exp(\boldsymbol{\delta}'\boldsymbol{z_i}) \qquad (3a)$$

$$\mu_i = \exp(\boldsymbol{\delta}'\boldsymbol{z_i}) \qquad (3b)$$

$$\theta_i = \theta\exp(\boldsymbol{\delta}'\boldsymbol{z_i}) \qquad (3c)$$

[13] Some regions can have CDD or HDD that are zero for some years, so taking a log isn't possible without dropping those observations.

Where $\boldsymbol{z_i}$ is a vector of industry group dummy variables for the ith plant based on which of the industry groups the plant belongs. While each of these modifications has a similar impact, they have a slightly different in interpretation. In the case of half normal and exponential, this approach amounts to heteroscedasticity with respect to industry group. In the case of truncated normal the truncation point is industry specific, but the variance remains the same. They all imply that the distribution of efficiency within a sector may vary, relative to the common estimated frontier demand function.

It is important to note that the "common frontier" includes industry level (6-digit NAICS) fixed effects so the frontier is industry specific but the elasticities are assumed to be common across industry group. The additional parameterization of the efficiency term above allows for the possibility that the distribution of efficiency varies across the six study industry groups.

This approach also does not consider any panel effects. It implicitly assumes that a manufacturing plant in the data is independent any prior (or subsequent year) performance. (Greene 2005) lays out the econometric approach for a "True Random Effects" and "True Fixed Effects" frontier specification. In practice, such models are very difficult to estimate when the number of observations are large compared to the number of time steps, which is the case here. Attempts to directly estimate these types of panel models failed to converge in every instance. However, since the time steps involved are quite long, every five years, the independence assumption is a bit easier to rationalize. In the next section an alternative to the direct estimation of panel effects is pursued.

Modelling plant heterogeneity in panel data

In a panel data context, it becomes possible to consider both unobserved heterogeneity and efficiency. Heterogeneity would be unmeasured differences in products that require more (or less) energy to manufacture or material choices. Efficiency includes both managerial and technology choices that result in lower energy use to achieve the desired production activity. One direct approach to account for plant heterogeneity would be to use detailed material and product codes to capture this using observable product mix. This has been done by (Boyd 2008, Boyd and Guo 2014, Boyd 2016) for some selected industries, but requires a large amount of knowledge regarding which specific material and product types are most relevant. Use of capital stock and material purchases may also account for these observable plant level differences, since some more energy intensive plants are likely to have less expensive material inputs since they may make, rather than buy, some intermediate product. Making an intermediate product instead of purchasing it is more likely to be both more energy intensive and more capital intensive as well. Even though we include capital stock and material purchases in the specification, additional methods to account for plant level heterogeneity may be desirable.

The desire to distinguish between efficiency and heterogeneity requires an extension of the SFA frame work to a panel-data setting. The standard treatment for plant level heterogeneity in panel data is to include either a plant specific fixed or random effect. Equation (4) represents the non-stochastic frontier implementation of plant level heterogeneity by the inclusion of ω_i, for the ith plant. ω_i, may be estimated by either a fixed or random effects estimator.

$$E_{i,t} = f\left(X_{i,t}: \theta\right) + \omega_i + \varepsilon_{i,t} \qquad (4)$$

The form of the model $f\left(X_{i,t}: \theta\right)$ would be the same as in equation (1). In the SFA approach the typical error term is hypothesized to be made up of two parts,

$$\varepsilon_{i,t} = u_{i,t} + v_{i,t} \qquad (5)$$

Where $u_{i,t}$ is a one-sided efficiency error term and $v_{i,t}$ is noise. (Greene 2005) shows that this extension of the SF framework is econometrically tractable via maximum likelihood estimation (MLE). This approach has been labeled Greene's true fixed effect (TFE) and true random effect (TRE) estimators. In the TRE model, the estimates of ω_i are the basis for an estimate of persistent efficiency and $u_{i,t}$ is time varying efficiency. (Filippini and Hunt 2011, Filippini and Hunt 2012) employ this approach on panels of US states and OECD countries, respectively. However, these models can be difficult to obtain convergence in the MLE when the number of time periods is relatively small (in this particular study 5 time periods) and the number of plants is relatively large. This was the same problem reported by (Boyd and Lee 2019) and is the case here as well. In this case the smallest sample size was 200 (*wet corn refining*), with most the largest 34,500 (*Other*). The panel is also unbalanced so each plant may not appear in the data for all five years.

An alternative approach is to estimate these error components in a two stage process (Kumbhakar, Lien et al. 2014). The next section describes the two stage process. The advantage here is that the convergence problems are ameliorated and heterogeneity can be treated in the first stage using a fixed effects model.

Two stage model for persistent and time varying efficiency

The plant level efficiency estimates are obtained by a two stage approach. The first stage uses a plant level either fixed or random effects estimator. Based on a Hausman test, fixed effects was employed. The general form for the estimate is shown above in equation (4). Where ω_i is the plant level fixed effect for the i[th] plant and $\varepsilon_{i,t}$ is Gaussian error. These two error components are not directly observable, but the residual of the regression, $E_{i,t} - f(X_{i,t}{:}\,\hat{\theta})$, can be decomposed into an estimate of the plant specific effect, $\widehat{\omega_i}$ that is constant over time for each plant and the time varying noise component, $\widehat{\varepsilon_{i,t}}$, based on the estimated parameters, $\hat{\theta}$.

$$\widehat{\omega_i} = E[\omega_i{:}\, E_{i,t} - f(X_{i,t}{:}\,\hat{\theta}), \hat{\theta}] \qquad (6a)$$

$$\widehat{\varepsilon_{i,t}} = E[\varepsilon_{i,t}{:}\, E_{i,t} - f(X_{i,t}{:}\,\hat{\theta}), \hat{\theta}] \qquad (6b)$$

The second stage is used to further extract efficiency estimates from the decomposed error terms using a stochastic frontier. The second stage applies a stochastic frontier regression to the plant level data derived from the fixed or random effects estimates, $\widehat{\omega_i}$ and $\widehat{\varepsilon_{i,t}}$, with no explanatory variables and a simple intercept term α.

$$\widehat{\omega_i} = \alpha + u_i^{per} + v_i \qquad (7a)$$

$$\widehat{\varepsilon_{i,t}} = \alpha + u_{i,t}^{tv} + v_{i,t} \qquad (7b)$$

The "usual" stochastic frontier model assumptions apply; u_i^{per} and $u_{i,t}^{tv}$ follow a one-sided distribution and v_i and $v_{i,t}$ are noise. In this study the exponential distribution was used for u_i^{per} and $u_{i,t}^{tv}$. We are not interested in the estimate, $\hat{\alpha}$, per se, but in the estimates of $\widehat{u_{i,t}^{tv}}$ and $\widehat{u_i^{per}}$ based on the residuals, $\widehat{\omega_i} - \hat{\alpha}$ and $\widehat{\varepsilon_{i,t}} - \hat{\alpha}$, from each regression. The standard JMLS (Jondrow, Materov et al. 1982) frontier estimates of $\widehat{u_{i,t}^{tv}}$ and $\widehat{u_i^{per}}$ are obtained from these two 2[nd] stage regressions . The exponent of these JMLS estimates represent time-varying (*tv*) and persistent (*per*) efficiency.

$$tv_{i,t} = \exp(\widehat{u_{i,t}{}^{tv}}) , \qquad (8a)$$

$$per_i = \exp(\widehat{u_i{}^{per}}) , \qquad (8b)$$

$$tot_{i,t} = tv_{i,t} \cdot per_i = \exp\left(\widehat{u_{i,t}{}^{tv}} + \widehat{u_i{}^{per}}\right) \qquad (8c).$$

Where $tot_{i,t}$ is the combined total efficiency estimate and is equal to $tv_{i,t} \cdot per_i$, since the two sources of efficiency are multiplicative as a result of the original equation being estimated in logs.

Frontier Efficiency Estimates

In addition to the energy data issues regarding fuel use in NAICS 311211 and 31131, initial analysis using the four IDM groupings found high levels of inefficiency in IDM-1 and IDM-4. While IDM-4, i.e. "other," might be expected to have a wider range of efficiency and lower average efficiency, estimates for average efficiency in IDM-1 were extremely low. Much of this inefficiency in IDM-1 was clustered in NAICS 311211; similarly, for NAICS 31131 in IDM-4. For this and the energy accounting issues identified regarding higher levels of coal and bio-mass based fuel versus natural gas use in these sectors, the IDM sectors were disaggregated for the purposes of the SFA analysis.

Table 2 presents the estimates of a common pooled frontier with 6-digit NAICS and year fixed effects. The parameter estimates of the frontiers all are significant and have the expected signs. All coefficients are fairly similar across the models using different efficiency distributions. Recall that since the model is estimated in logs, the coefficients are interpreted as elasticities, except for CDD, HDD, and GERATIO. CDD and HDD have the expected signs and are significant. Own-price elasticities range from -0.7 to -0.9, with cross price elasticities, i.e. the coefficient of the electric price in the fuel equation and vice-versa, that are rather small but significant in the neighborhood of 0.04 to 0.1. This suggests that electricity and fuels are used to meet rather specific end uses, e.g. motor drive, and there is limited substitution between electricity and fuel, i.e. natural gas, coal, petroleum etc. The model does not estimate possible substitution between fuels, e.g. coal for natural gas. The model exogenously controls for the decision to employ CHP by including the GERATIO variable, defined as the ratio of on-site generated power to net on-site consumption. GERATIO is significant and has the expected signs, positive for fuel and negative for electricity. Including GERATIO implies that these cross price elasticities do not include substitution due to the use of CHP. Some type of two-stage model would be needed to account for the decision to install cogeneration capability. That is left for future study.

Mean electricity efficiency estimates are also similar across models, between 0.7 and 0.75. For fuel, one model estimate is of efficiency is 0.65 while the other two are much higher at 0.9. This modeling approach is designed to test if the efficiency distributions differ across industry types, compared to a common frontier. The parameterization of the efficiency distribution measures the differences in the variances of five different industry groups compared to the reference industry group (IG-1: *NAICS 3112 less 311211*). These estimates are generally significant, implying that most industries have a different variance for the efficiency distribution from IG-1. It is rather difficult to directly interpret these coefficients, so a more intuitive comparison of the differences in efficiency are presented graphically below.

Table 2 Common frontier estimates under alternative frontier specifications

		Exponential		half normal		truncated normal	
		electricity	fuel	electricity	fuel	electricity	fuel
Frontier	Lny	0.400***	0.354***	0.392***	0.334***	0.391***	0.349***
	Lnk	0.200***	0.193***	0.198***	0.193***	0.198***	0.196***
	Lnl	0.278***	0.320***	0.283***	0.336***	0.283***	0.320***
	Lnm	0.163***	0.131***	0.165***	0.137***	0.166***	0.133***
	lnPrE	-0.729***	0.0699***	-0.729***	0.0581***	-0.726***	0.0463***
	lnPrF	0.0519***	-0.857***	0.0981***	-0.767***	0.0976***	-0.757***
	GERATIO	-1.383***	0.998***	-1.315***	0.859***	-1.379***	0.979***
	CDD	3.80e-05***		3.96e-05***		3.98e-05***	
	HDD		2.01e-05***		1.91e-05***		1.79e-05***
	Constant	-1.355***	3.655***	-1.451***	3.061***	-1.442***	3.440***
$\sigma_{u,i}^2$	2.indgroup	-2.936***	-2.552***	-1.966***	-0.374***		
	3.indgroup	-2.219***	-2.287***	-1.298***	-0.446***		
	4.indgroup	-2.283***	-7.108***	-1.400***	-1.201***		
	5.indgroup	-1.462***	-3.649**	-2.197	0.142		
	6.indgroup	-0.640***	-0.645**	-0.282	0.15		
Mu	2.indgroup					-1.040***	-2.643***
	3.indgroup					-0.478***	-2.283***
	4.indgroup					-0.634***	-37.81***
	5.indgroup					0.55	-5.9
	6.indgroup					0.994***	8.001
	Observations	48500	48500	48500	48500	48500	48500
	E[sigma_u]	0.347	0.142	0.513	0.637	0.632	1.008
	sigma_v	0.545	0.894	0.564	0.842	0.564	0.895
	Mean Efficiency	0.753	0.9	0.698	0.647	0.738	0.893

*** $p<0.01$ ** $p<0.05$ * $p<0.1$

While the distributional assumptions used to obtain the SFA estimates are illustrated in figure 1, the empirical distributions using non-parametric representations (kernel density) of the efficiency distributions by energy type, industry group, under the three different frontier distributions assumptions are shown in Figures 2-7[14]. *Assuming a common frontier for the continuous covariates*, the industry group level distributions are very sensitive to the assumptions regarding the parametric representation of efficiency in the stochastic frontier estimates. It is important to point out the "common frontier" estimates include industry fixed effects (6 digit NAICS), so does control for industry specific conditions at a detailed level. Given this observation, it would be natural to try to relax the common frontier assumptions for the covariates and estimate individual frontiers for each industry group. When that was attempted the frontiers rarely converged. This same problem was encountered by (Boyd and Lee 2020) for the chemical industry. To address the desire to estimate sector specific frontiers the two stage method proposed by (Kumbhakar, Lien et al. 2014) and used by (Boyd and Lee 2020) is presented next.

[14] Some errors appear to have occurred in generating the graphs in figures 2 and 3 during the clearance process, these will be fixed once access to the RDC has resumed.

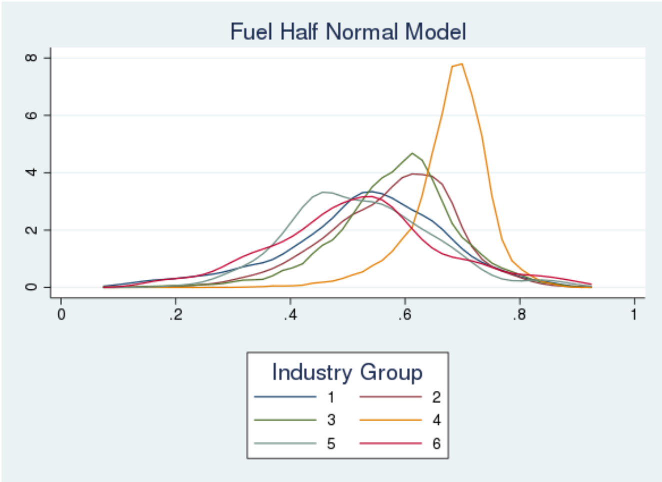

Figure 2 Kernel Density Estimates of Plant Level Efficiency by Energy, Industry Group, and Frontier type

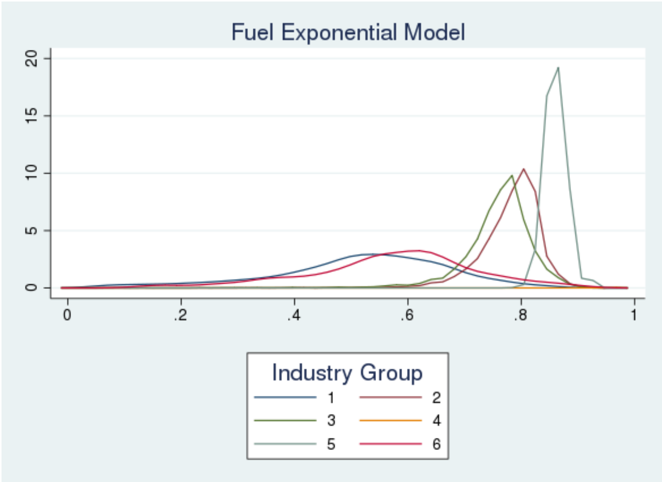

Figure 3 Kernel Density Estimates of Plant Level Efficiency by Energy, Industry Group, and Frontier type

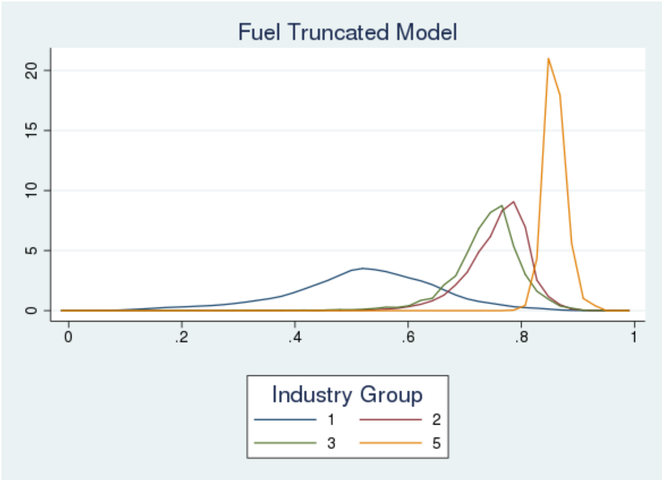

Figure 4 Kernel Density Estimates of Plant Level Efficiency by Energy, Industry Group, and Frontier type (convergence issues prevented kernel density estimates for IG-4 and IG-6 for this model)

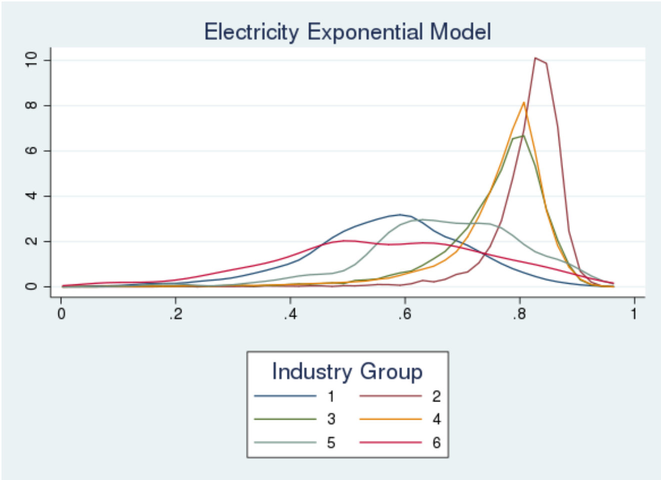

Figure 5 Kernel Density Estimates of Plant Level Efficiency by Energy, Industry Group, and Frontier type

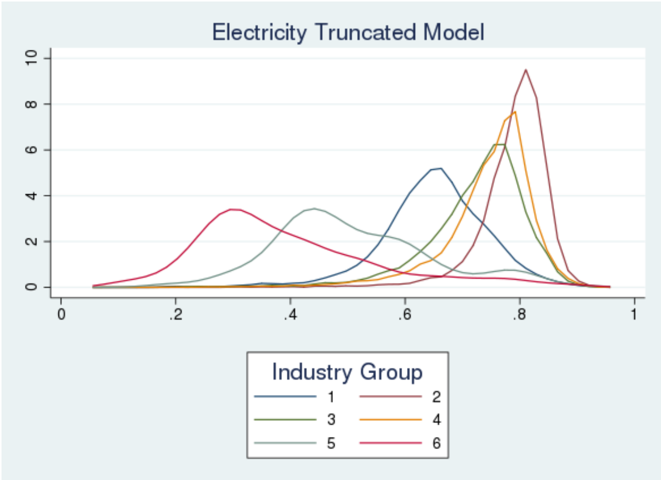

Figure 6 Kernel Density Estimates of Plant Level Efficiency by Energy, Industry Group, and Frontier type

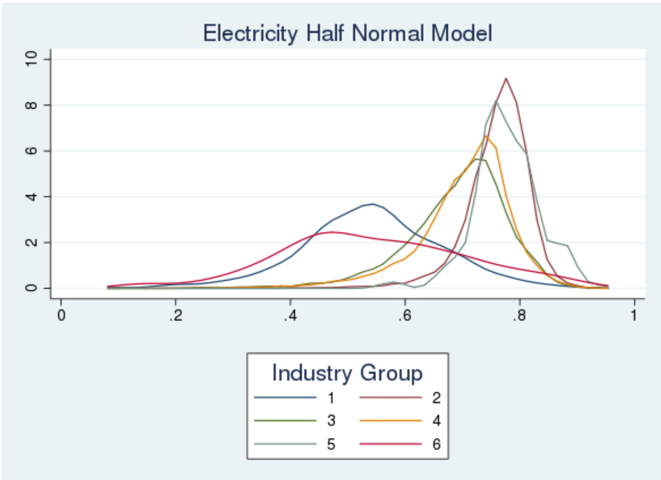

Figure 7 Kernel Density Estimates of Plant Level Efficiency by Energy, Industry Group, and Frontier type

Tables 3 and 4 provides the estimates from the two stage models for electricity and fuels restrictively[15]. The estimates for the industry specific demand frontiers using the two stage estimator show that not only does efficiency vary across industry group, so do the demand frontier parameter estimates. Output elasticities vary quite a bit across sectors. IG-3 and IG-4 are the most similar to the pooled model in table 2. Own price elasticities are the same or lower, in IG-1 in particular. Material elasticities are much higher in IG-5 and IG-6. These are the two sectors that were disaggregated. They are also two of the most "upstream" food processing sectors that either turn corn into starch and sugars (IG-5, NAICS 311211) or turn cane and beet into sugar (IG-6, NAICS 31131). The importance of the volume of primary materials as an energy driver is therefore not surprising. The use of plant fixed effects may explain the difficulty in getting significant CDD and HDD coefficients. Most of the variation in CDD and HDD are locational (i.e. climate) and would be absorbed by the fixed effects. The remaining time series variations might make reliable estimates of "weather vs climate" more difficult.

Focusing on the rows labeled permanent, transient, and total efficiency the value range from 0.639 to 0.982 for electricity and 0.55 to 0.97 for fuel. Total efficiency is always lower since it is the combined results of permanent and transient efficiency. One can interpret this as having average levels of efficiency are broken down into transient and permanent. When we focus on total efficiency we see that the range is 55% to 88%, with average fuel efficiency either about the same or lower than for electricity, with the exception of IG-6. IG-4 has the lowest level if total efficiency; it is also the most diverse sector. Permanent efficiency tends to be higher across industry and fuel type, with the exception of IG-4. This low level of permanent efficiency contributes to the lower level of total efficiency compared to the other sectors. Transient efficiency ranges from 80%-90%.

Table 3 Two stage electricity demand frontier by industry group

VARIABLES	IG-1	IG-2	IG-3	IG-4	IG-5	IG-6
Lny	0.186***	0.363***	0.464***	0.311***	0.21	0.253
Lnk	0.0825**	0.115***	0.107***	0.0778***	0.0923	0.0754**
Lnm	0.194***	0.101**	0.0207	0.193***	0.383**	0.376***
Lnl	0.352***	0.237***	0.233***	0.265***	-0.209	0.491***
lnPrE	-0.588***	-0.605***	-0.721***	-0.790***	-0.754**	-0.879***
lnPrF	0.0237	0.0402	0.00915	0.0707**	0.551**	0.690**
GERATIO	-1.718***	-0.246	-0.348	-0.867***	-0.694	-0.339
CDD[16]	0.0168	0.000764	-0.181***	-.00581	0.131	-0.665
Constant	3.483***	1.916***	1.804***	0.181	3.066	-1.795
Observations	2200	4400	6900	34500	200	350
Number of Plants	850	1700	2900	17500	70	100
R-squared	0.437	0.537	0.457	0.478	0.367	0.51
Permanent Efficiency	0.871	0.982	0.979	0.738	0.862	0.833
Transient Efficiency	0.888	0.899	0.871	0.866	0.862	0.873
Total Efficiency	0.773	0.883	0.853	0.639	0.745	0.728

*** p<0.01 **p<0.05 * p<0.1

[15] Some of the frontier estimates for permanent efficiency fail to converge after 100 iterations.
[16] Coefficients on degree days are scaled to be per thousand CDD and HDD.

Table 4 Two stage fuel demand frontier by industry group

VARIABLES	IG1	IG-2	IG-3	IG-4	IG-5	IG-6
Lny	0.209**	0.415***	0.435***	0.299***	0.143	0.0851
Lnk	0.102***	0.052	0.0896***	0.0163	0.357	0.0818
Lnm	0.145*	0.113**	0.0751***	0.185***	0.448**	0.445*
Lnl	0.414***	0.188***	0.195***	0.303***	-0.162	0.181
lnPrE	0.0249	0.140***	-0.0239	0.173***	0.0836	-0.589***
lnPrF	-0.645***	-0.728***	-0.878***	-0.745***	0.106	-0.327
GERATIO	0.386	0.0398	0.104	0.262	0.478	-0.0386
HDD[16]	0.0274	-0.00894	0.181***	0.00534	0.536	0.065
Constant	4.615***	4.863***	3.799***	3.645***	0.995	3.144
Observations	2200	4400	6900	34500	200	350
Number of Plants	850	1700	2900	17500	70	100
R-squared	0.169	0.241	0.21	0.253	0.315	0.392
Permanent Efficiency	0.941	0.856	0.970	0.668	0.845	0.956
Transient Efficiency	0.830	0.805	0.815	0.828	0.837	0.823
Total Efficiency	0.781	0.689	0.790	0.553	0.706	0.787

*** p<0.01 **p<0.05 * p<0.1

One goal of this study is to estimate efficiency levels for the four NEMS food industry groups. IG-2 and IG-3 are the same as IDM-2 and IDM-3, but IDM-1 and IDM 4 were disaggregated. Using the models based on the study industry groups, estimates for IDM-1 and IDM-4 average efficiency can be constructed. Table 5 presents those estimates using both the simple average and the plant energy weighted average for total efficiency. The values in grey correspond to those reported in table 3. Using the energy weighted average results in slightly lower efficiencies. It is not a given that the weighted averages would be lower, but it depends on the relationship between the plant size distribution relative to the plant efficiency distribution. While the model does control for size, it appears that size in not completely random relative to estimated plant energy efficiency. The differences between weighted and un-weighted are small in three out of the eight estimates. The largest difference is in IDM-4; 11% and 12% for fuel and electricity respectively.

Table 5 Estimates of total efficiency from the two stage model by NEMS IDM groups

NEMS IDM	1	2	3	4
mean total electricity efficiency	78%	88%	85%	64%
plant weighted total electricity efficiency	77%	87%	84%	52%
mean total fuel efficiency	80%	69%	79%	56%
plant weighted total fuel efficiency	77%	64%	74%	45%

Efficiency of new firms

It is commonly assumed that new plants entering an industry are likely to employ the newest and most advanced production technology. To empirically assess this vis-à-vis energy, we compute the energy efficiency during the year that a plant enters the industry and compare it to plants that existed in the prior time period. We conduct the analysis using the NEMS industry groups, rather than the disaggregated sectors. This is done to allow these estimates to be used directly in the NEMS model. In terms of the data, entry is considered to be when the plant first appears in the CMF data. This means that they entered sometime in the five-year window proceeding the first year they appear in the data.

This efficiency measure is year specific since year fixed effects are used. We pool all new entrants and compare them to the incumbents. It should be noted that a plant can be an entrant, then later an incumbent. In fact, all plants that enter after 1992 but before 2012 are also counted as incumbents five years after they entered. The share of total value of shipments represented by new plants ranges from 7% to 11%. We use a simple t-test for the mean level of energy efficiency between the two groups. Table 6 shows that new plants have significantly higher efficiency, but the values are not large, ranging from 0.4% to 5.7%. This is consistent with other results (Boyd and Lee 2019, Boyd and Lee 2020). Table 7 and 8 show the breakdown for permanent and transient efficiency. New plants are also statistically more efficient when viewed by the efficiency components, with the one exception; permanent fuel use in IDM- 1. With the exception of IDM-4, transient efficiency differences dominate the differences.

Table 6 New vs Existing Plant Efficiency by NEMS Industry Group – Total Efficiency

NEMS industry group	IDM-1	IDM-2	IDM-3	IDM-4
Electricity				
Existing	77.7%	88.3%	85.1%	62.2%
New	78.4%	88.6%	85.8%	67.3%
Difference	0.8%	0.4%	0.7%	5.1%
t-test	-2.595	-2.029	-4.4	-45.04
Fuel				
Existing	79.5%	68.6%	79.0%	53.5%
New	80.6%	70.4%	79.8%	59.2%
Difference	1.2%	1.7%	0.8%	5.7%
t-test	-3.402	-5.39	-4.236	-46.44

Table 7 New vs Existing Plant Efficiency by NEMS Industry Group – Permanent Efficiency

NEMS industry group	IDM-1	IDM-2	IDM-3	IDM-4
Electricity				
Existing	87.8%	98.3%	97.9%	72.2%
New	88.2%	98.3%	98.0%	76.9%
Difference	0.3%	0.0%	0.0%	4.7%
t-test	-1.785	-8.81	-20.95	-40.57
Fuel				
Existing	95.9%	85.5%	97.1%	65.0%
New	95.9%	86.2%	97.1%	70.9%
Difference	0.0%	0.7%	0.0%	5.9%
t-test	-0.8427	-7.161	-14.85	-43.82

Table 8 New vs Existing Plant Efficiency by NEMS Industry Group – Transient Efficiency

NEMS industry group	IDM-1	IDM-2	IDM-3	IDM-4
Electricity				
Existing	88.4%	89.8%	86.9%	86.2%
New	89.0%	90.2%	87.6%	87.5%
Difference	0.6%	0.3%	0.7%	1.3%
t-test	-1.779	-1.933	-4.189	-18.59
Fuel				
Existing	82.9%	80.3%	81.3%	82.3%
New	84.1%	81.6%	82.2%	83.6%
Difference	1.2%	1.3%	0.8%	1.2%
t-test	-3.384	-3.629	-4.087	-15.23

Alternative estimates of potential efficiency improvement

Figures 1-6 illustrate how the efficiency distributions vary across different industries and specific energy types. It is also clear that these estimates imply that few plants achieve 100% efficiency. This is in part, an artifact of the SFA approach and in particular the method used to decompose the residual into an efficiency component and the error component (Jondrow, Materov et al. 1982). This decomposition can place some, or even many, plants close to the frontier, i.e. 100%. What is clear from figures 2-7, that is not the case with these empirical estimates.

The fact that few plants achieve anything close to 100% efficiency has important implications to the interpretation of these results. One could interpret this as empirical evidence of both a *maximum level of efficiency* based on the frontier vs a *practical level of efficiency* based on what is more commonly observed in the data. The estimated frontier can be interpreted as technically feasible, i.e. estimated from the data using the assumptions of the SFA approach, but 100% efficiency may not be achieved very often. This means that we will want to have alternative reference points for assessing efficiency.

The energy efficiency literature has long distinguished between different notions of efficiency. (Jaffe and Stavins 1994) introduce the notion of three types of "potential"; *economic*, *technical*, and *hypothetical*. The different levels of efficiency in their taxonomy are respectively related to market failures for efficiency, elimination of uncertainty/inertia/heterogeneity, and market failures generally. In this study efficiency is estimated based on observed behavior. To the extent that few, if any, firms achieve the 100% efficient frontier one might consider the frontier estimate to be similar to the hypothetical (maximum) level of performance. The other two definitions in their taxonomy, *economic* and *technical*, relate to market and non-market barriers. This approach is not able to identify the sources of efficiency. However, we can estimate energy reductions that could be achieved by moving low performing plants to achieve energy use similar to those of higher performing plants. We propose to define *practical efficiency* as based on some level of observed practice, where *observed practice is measured by a selected percentile of the efficiency distribution.* Using this approach to measuring practical efficiency, we need not assume that firms can achieve 100% efficiency, particularly when it is rarely achieved in practice.

To do this, we chose a percentile that is considered the target (desired) level of practical efficiency. For example, the 75th percentile or above is considered by the U.S. EPA Energy Star program to be "efficient". The frontier estimates reported above can provide an estimate of the level of efficiency that is at the 75th percentile of the efficiency distribution. From this perspective we define "efficient" as performance that is at least as good as 3/4th of all plants in the industry. For illustration, assume that for a particular industry and energy type the empirical level of efficiency at the 75th percentile, relative to the frontier of 100%, is 60%. We can compute the energy use for all plants whose efficiency is below 60%, assuming that they raise their efficiency to 60%. If a plant is 40% efficient then it can lower its energy use by 1/3rd to achieve 60% efficiency.

For calculations of "practical efficiency" the NEMS industry groupings are used, but with the disaggregated study grouping results. Table 9 summarizes the savings for 95th percentiles, using the approach described above. The simulation is a percentage of total energy use that would result if all plants below the target level of the cumulative efficiency distribution were to improve and those above would remain the same. This can be interpreted as targeting a "practical" level of efficiency instead of assuming a hypothetical level of 100%. Weighted plant efficiencies are generally lower, but this is not universal. In particular, IG-1 has higher weighted efficiency in several of the model estimates. As was observed above, the estimated efficiency distributions are sensitive to the model specifications. If we take a simplistic view and average across estimates the unweighted efficiencies range from 73%-93% depending on energy and industry types. Weighted efficiency ranges from 71% to 85%, but the low and high ends of those ranges do not coincide.

This exercise can use different target levels of practical efficiency by choosing different percentiles. We choose the 50th and 75th. One could think of these as the result of policy that targets "low handing fruit" or setting a performance based efficiency targets like those used in voluntary programs such as Energy Star[17].

Tables 10 and 11 show estimates of practical efficiency at the 75th and 50th percentiles, respectively. As we lower the target percentile level in the efficiency distribution the "energy savings" estimates of practical efficiency decline, by definition. These estimates are all higher than the mean efficiency estimates presented in table 9, since those numbers implicitly assume a higher level of target practical efficiency, i.e. the 95th percentile. Table 10 shows the estimates by energy type, industry, and model based on the 75th percentile. Focusing on the model averaged, plant weighted simulated efficiencies, the difference between practical efficiency defined by the 95th vs the 75th is only about 5%, i.e. from ~80% at the 95th percentile and ~85% at the 75th percentile. While there are potential energy savings of 20%, assuming a more stringent level of target practical efficiency of the top 5%, there is still 15% potential energy savings at a "less stringent" target.

Table 10 shows the calculation of practical efficiency defined at the 50th percentile. Again focusing on the energy/industry plant weighed estimates, if the target level of efficiency is improving the lower half of the efficiency distribution, i.e. "low hanging fruit", the level of savings potential drops, but only by another 2.4% from the savings potential at the 75th. A large part of the efficiency gap is in the lower half of the distribution. Based on a target level of efficiency using the 50th percentile, i.e. assuming that the worst performing plants simply achieve the median level of performance, efficiency estimates

[17] The Energy Star program uses the upper quartile of performance as the level that can be certified as Energy Star efficient level. The 75th percentile or above is used in the Energy Star Commercial and Industrial Buildings program to offer manufacturing plant and commercial building Energy Star certification (see https://www.energystar.gov/buildings/about-us)

ranges from 75% to 95%. This translates to from a 5% to 25% saving potential, depending on the energy type and industry. The simple average is 87%

Table 9 Weighted and unweighted efficiency estimates based on the 95th percentile of the cumulative distributions, by NEMS Industry group, type of energy, and SFA assumptions

NEMS Industry Grouping	IDM-1	IDM-2	IDM-3	IDM-4
Two Stage Regression model				
mean total electricity efficiency	92%	96%	94%	85%
plant weighted total electricity efficiency	91%	95%	93%	69%
mean total fuel efficiency	93%	90%	92%	80%
plant weighted total fuel efficiency	90%	83%	86%	63%
Exponential stochastic frontier model				
mean total electricity efficiency	72%	92%	87%	88%
plant weighted total electricity efficiency	85%	77%	79%	79%
mean total fuel efficiency	64%	92%	90%	99%
plant weighted total fuel efficiency	86%	87%	87%	80%
Half Normal stochastic frontier model				
mean total electricity efficiency	70%	91%	85%	86%
plant weighted total electricity efficiency	82%	78%	80%	80%
mean total fuel efficiency	71%	78%	78%	88%
plant weighted total fuel efficiency	70%	70%	69%	67%
Truncated stochastic frontier model				
mean total electricity efficiency	80%	92%	86%	88%
plant weighted total electricity efficiency	83%	77%	79%	78%
mean total fuel efficiency	64%	91%	89%	99%
plant weighted total fuel efficiency	81%	82%	82%	75%
Average across models				
mean total electricity efficiency	79%	93%	88%	87%
plant weighted total electricity efficiency	85%	82%	83%	77%
mean total fuel efficiency	73%	88%	87%	92%
plant weighted total fuel efficiency	82%	80%	81%	71%

Table 10 Weighted and unweighted efficiency estimates based on the 75th percentile of the cumulative distributions, by NEMS Industry group, type of energy, and SFA assumptions

NEMS Industry Grouping	IDM-1	IDM-2	IDM-3	IDM-4
Two Stage Regression model				
mean total electricity efficiency	98%	98%	98%	93%
plant weighted total electricity efficiency	95%	97%	97%	78%
mean total fuel efficiency	97%	96%	97%	91%
plant weighted total fuel efficiency	95%	90%	92%	71%
Exponential stochastic frontier model				
mean total electricity efficiency	83%	95%	92%	92%
plant weighted total electricity efficiency	91%	80%	83%	83%
mean total fuel efficiency	79%	95%	95%	100%
plant weighted total fuel efficiency	93%	88%	89%	80%
Half Normal stochastic frontier model				
mean total electricity efficiency	83%	95%	91%	92%
plant weighted total electricity efficiency	93%	82%	86%	86%
mean total fuel efficiency	82%	86%	88%	93%
plant weighted total fuel efficiency	81%	78%	78%	72%
Truncated stochastic frontier model				
mean total electricity efficiency	89%	95%	92%	92%
plant weighted total electricity efficiency	89%	80%	84%	83%
mean total fuel efficiency	82%	95%	94%	99%
plant weighted total fuel efficiency	87%	83%	84%	75%
Average across models				
mean total electricity efficiency	88%	96%	93%	91%
plant weighted total electricity efficiency	92%	84%	87%	81%
mean total fuel efficiency	85%	92%	93%	94%
plant weighted total fuel efficiency	89%	84%	85%	73%

Table 11 Weighted and unweighted efficiency estimates based on the 50th percentile of the cumulative distributions, by NEMS Industry group, type of energy, and SFA assumptions

NEMS Industry Grouping	IDM-1	IDM-2	IDM-3	IDM-4
Two Stage Regression model				
mean total electricity efficiency	98%	98%	98%	93%
plant weighted total electricity efficiency	95%	97%	97%	78%
mean total fuel efficiency	97%	96%	97%	91%
plant weighted total fuel efficiency	95%	90%	92%	71%
Exponential stochastic frontier model				
mean total electricity efficiency	90%	97%	94%	94%
plant weighted total electricity efficiency	94%	82%	86%	85%
mean total fuel efficiency	87%	97%	97%	100%
plant weighted total fuel efficiency	95%	89%	90%	80%
Half Normal stochastic frontier model				
mean total electricity efficiency	91%	97%	94%	95%
plant weighted total electricity efficiency	96%	85%	90%	89%
mean total fuel efficiency	89%	92%	93%	96%
plant weighted total fuel efficiency	88%	84%	83%	75%
Truncated stochastic frontier model				
mean total electricity efficiency	93%	97%	95%	95%
plant weighted total electricity efficiency	92%	82%	87%	86%
mean total fuel efficiency	90%	97%	96%	99%
plant weighted total fuel efficiency	88%	84%	85%	75%
Average across models				
mean total electricity efficiency	93%	97%	95%	94%
plant weighted total electricity efficiency	94%	86%	90%	84%
mean total fuel efficiency	91%	95%	96%	96%
plant weighted total fuel efficiency	92%	87%	87%	75%

Summary

This paper presents analysis of the energy efficiency of the food processing industry based on non-public plant level data from the US Census Bureau. Estimates using four NEMS industry groupings and a disaggregation of two of those sectors, for a total of six sub-sectors of the food processing industry, were developed under several different SFA methodologies. The SFA models are based on an energy demand specification that control for plant characteristics such as output, capital labor, and materials in addition to energy prices. Controls for energy related climate (annual HDD and CDD), detailed industry codes (6-digit NAICS), year, and combined heat and power were also included. Own price elasticities range from -0.6 to -0.9, with little evidence of significant cross price substitution. As mentioned above, CHP is controlled for separately, so these estimates of substitution are “direct” and do not include the decision to employ this technology. This means that the estimates would therefore be lower than if we didn’t control for this as an exogenous effect.

The estimated energy efficiency distributions are sensitive to the choice of frontier modeling assumptions and also reveal that few plants achieve anything near 100% efficiency. Using a two stage approach, estimates of persistent and transient efficiency are also obtained. Average transient efficiency was generally, but not universally, the larger contributor to total efficiency. New plant, on average, were statistically more efficient than existing plants for both permanent and transient efficiency, but the values are not large, ranging from a low of 0.4% to a high of 5.7%, depending on the industry group and energy type. NEMS industry group 4, "Other" exhibited the largest difference between new and continuing plants.

Since few plants achieve the maximum level of 100% efficiency, the average efficiency is not a useful estimate of what may be practically achievable. We introduce the notion of practical efficiency, based on the full distribution and a target percentile. If we simulate efficiency improvements, i.e. the "efficiency gap," by setting the target level of efficiency based on a percentile on the cumulative efficiency distribution the estimate of potential reduction in energy use is by definition smaller, but empirically more reasonable. A large portion of energy saving from the estimated energy efficiency gap can be a achieved by focusing on "low hanging fruit," defined as bringing the bottom half of the distribution up to the median, resulting in an average 13% reduction. If the target level of efficiency is the upper quartile, i.e. the same as used by Energy Star for certification of industrial plants (75th percentile), an additional 2.5% savings is achievable. Setting the practical level of energy efficiency at the 95th percentile a 20% reduction is achievable; an additional 4.5% above the Energy Star level of the 75th percentile and 7% more than the "low hanging fruit" of the 50th percentile.

References

Allcott, H. and M. Greenstone (2012). "Is There an Energy Efficiency Gap?" The Journal of Economic Perspectives **26**(1): 3-28.

Bardazzi, R., F. Oropallo and M. G. Pazienza (2015). "Do manufacturing firms react to energy prices? Evidence from Italy." Energy Economics **49**: 168-181.

Blumstein, C. and S. Stoft (1995). "Technical Efficiency, production functions, and conservation supply curves." Energy Policy **23**(9): 765-768.

Boyd, G. and G. Zhang (2013). "Measuring improvement in energy efficiency of the US cement industry with the ENERGY STAR Energy Performance Indicator." Energy Efficiency **6**(1): 105-116.

Boyd, G. A. (2008). "Estimating Plant Level Manufacturing Energy Efficiency with Stochastic Frontier Regression." The Energy Journal **29**(2): 23-44.

Boyd, G. A. (2016). "Comparing the statistical distributions of energy efficiency in manufacturing: meta-analysis of 24 Case studies to develop industry-specific energy performance indicators (EPI)." Energy Efficiency: 1-22.

Boyd, G. A. and E. M. Curtis (2014). "Evidence of an "Energy-Management Gap" in U.S. manufacturing: Spillovers from firm management practices to energy efficiency." Journal of Environmental Economics and Management **68**(3): 463-479.

Boyd, G. A. and Y. Guo (2014). "An Energy Performance Indicator for integrated paper and paperboard mills: A new statistical model helps mills set energy efficiency targets." Paper360 **March/April**: 26-28.

Boyd, G. A. and J. M. Lee (2019). "Measuring plant level energy efficiency and technical change in the U.S. metal-based durable manufacturing sector using stochastic frontier analysis." Energy Economics **81**: 159-174.

Boyd, G. A. and J. M. Lee (2020). "Relative Effectiveness of Energy Efficiency Programs versus Market Based Climate Policies in the Chemical Industry." The Energy Journal **41**(3): 22.

Caves, D., L. Christensen and W. E. Diewert (1982). "The Economic Theory of Index Numbers and Measurement of Input, Output, and Productivity." Econometrica **50**(6): 1393-1414.

Davis, S. J., C. Grim, J. Haltiwanger and M. Streitwieser (2012). "Electricity Unit Value Prices and Purchase Quantities: U.S. Manufacturing Plants, 1963–2000." Review of Economics and Statistics **95**(4): 1150-1165.

Filippini, M. and L. C. Hunt (2011). "Energy Demand and Energy Efficiency in the OECD Countries: A Stochastic Demand Frontier Approach." The Energy Journal **32**(2): 59-80.

Filippini, M. and L. C. Hunt (2012). "US residential energy demand and energy efficiency: A stochastic demand frontier approach." Energy Economics **34**(5): 1484-1491.

Filippini, M. and L. C. Hunt (2015). "Measurement of energy efficiency based on economic foundations." Energy Economics **52, Supplement 1**: S5-S16.

Filippini, M. and L. C. Hunt (2016). "Measuring persistent and transient energy efficiency in the US." Energy Efficiency **9**(3): 663-675.

Foster, L., C. Grim and J. Haltiwanger (2016). Reallocation In The Great Recession: Cleansing Or Not? Working Paper, Center for Economic Studies, U.S. Census Bureau.

Gerarden, T. D., R. G. Newell and R. N. Stavins (2015). "Assessing the Energy-Efficiency Gap." National Bureau of Economic Research Working Paper Series **No. 20904**.

Greene, W. (2005). "Fixed and Random Effects in Stochastic Frontier Models." Journal of Productivity Analysis **23**(1): 7-32.

Howarth, R. B. and B. Andersson (1993). "Market Barriers to Energy Efficiency." Energy Economics(October): 262-272.

Huntington, H. (1995). "Been Top Down So Long It Looks Like Bottom Up." Energy.
Jaffe, A. B. and R. N. Stavins (1994). "The energy-efficiency gap What does it mean?" Energy Policy **22**(10): 804-810.

Jondrow, J., I. Materov, K. Lovell and P. Schmidt (1982). "On the Estimation of Technical Inefficiency in the Stochastic Frontier Production Function Model." Journal of Econometrics **19**(2/3): 233-238.

Kumbhakar, S. C., G. Lien and J. B. Hardaker (2014). "Technical efficiency in competing panel data models: a study of Norwegian grain farming." Journal of Productivity Analysis **41**(2): 321-337.

Rosenfeld, A., C. Atkinson, J. Koomey, A. Meier, R. J. Mowris and L. Price (1993). "Conserved Energy Supply Curves for U.S. Buildings " Contemporary Economic Policy **11**(1): 45-68.

Stoft, S. E. (1995). "The Economics of Conserved-Energy “Supply” Curves." *The Energy Journal* **16**(4): 109-137.

Worrell, E., S. Ramesohl and G. Boyd (2004). "ADVANCES IN ENERGY FORECASTING MODELS BASED ON ENGINEERING ECONOMICS." Annual Review of Environment and Resources **29**(1): 345-381.