Online Appendix for "Efficiency and Redistribution in Environmental Policy: An Equilibrium Analysis of Agricultural Supply Chains"

Tomás Domínguez-Iino*

Appendix A. Theory

A.1 Land use choice model

A.1.1 Specification of payoffs and idiosyncratic shocks

In each county i there is a continuum of fields, each indexed by ω. Each field ω is owned by a landowner, who chooses a land use from a discrete choice set consisting of a natural-use option $\mathcal{N}$ and a nest of agricultural commodities $\mathcal{C}$. We denote $r_i^c(\omega)$ as the landowner's return to commodity c, measured in dollars per unit of land, and $L_i^c(\omega)$ as the field's size (i.e., hectares). Returns are decomposed into the average return of county i, r_i^c, and a field ω idiosyncratic shock $\varepsilon_i^c(\omega)$,

$$r_i^c(\omega) = r_i^c \exp(\varepsilon_i^c(\omega)).$$

Let $A_i^{\mathcal{N}}(\omega)$ denote the payoff per unit of land when left in its natural state, which is decomposed into a county i average $A_i^{\mathcal{N}}$ and a field ω-specific idiosyncratic shock $\varepsilon_i^{\mathcal{N}}(\omega)$. Finally, let $r_i^k(\omega)$ denote the landowner's return per physical unit of land allocated to choice k,

$$r_i^k(\omega) = \begin{cases} r_i^c \exp(\varepsilon_i^c(\omega)) & \text{if } k = c \in \mathcal{C}, \\ A_i^{\mathcal{N}} \exp(\varepsilon_i^{\mathcal{N}}(\omega)) & \text{if } k = \mathcal{N}, \end{cases}$$

Since returns $r_i^k(\omega)$ are linear in land it is optimal to allocate the entire field to a single use (this is why the choice problem is discrete). Next, as in any discrete choice problem, the probability of making a specific choice k is invariant to monotonic transformations of payoffs. Therefore, consider the log returns, denoted $\tilde{r}_i^k(\omega)$,

$$\tilde{r}_i^k(\omega) = \begin{cases} \ln\left(r_i^c\right) + \varepsilon_i^c(\omega) & \text{if } k = c \in \mathcal{C}, \\ \ln\left(A_i^{\mathcal{N}}\right) + \varepsilon_i^{\mathcal{N}}(\omega) & \text{if } k = \mathcal{N}. \end{cases}$$

By having the field ω-idiosyncratic component of payoffs in additive form we can now use standard discrete choice results from industrial organization, which typically use additive type I EV errors instead of multiplicative type II EV (Fréchet) errors. Each $\varepsilon_i^k(\omega)$ is distributed type I extreme value with a dispersion parameter σ that determines the variance of the shocks across fields: the standard deviation of $\varepsilon_i^k(\omega)$ is $\sigma\frac{\pi}{\sqrt{6}}$. Because the scale of the payoffs is irrelevant for the landowner's discrete decision, we can re-scale them by any positive factor. Let θ denote the

*University of Chicago Booth School of Business. E-mail: tomasdi@uchicago.edu.

re-scaling factor, so that the re-scaled payoffs are $\tilde{r}_i^k(\omega)^* \equiv \theta \tilde{r}_i^k(\omega)$,

$$\tilde{r}_i^k(\omega)^* = \begin{cases} \theta \ln\left(r_i^c A_i^c\right) + \varepsilon_i^c(\omega)^* & \text{if } k = c \in \mathcal{C}, \\ \theta \ln\left(A_i^{\mathcal{N}}\right) + \varepsilon_i^{\mathcal{N}}(\omega)^* & \text{if } k = \mathcal{N}. \end{cases}$$

I follow the convention to set the re-scaling factor equal to the inverse of the type I EV dispersion parameter, i.e., $\theta = \sigma^{-1}$ (Train, 2009), which normalizes the standard deviation of the re-scaled error $\varepsilon_i^k(\omega)^* \equiv \theta \varepsilon_i^k(\omega)$ to $\frac{\pi}{\sqrt{6}}$. For this reason, the parameter θ governs the dispersion of the idiosyncratic shocks across fields (with higher values of θ corresponding to lower dispersion). Next, we turn to the assumptions that generate correlation between the $\varepsilon_i^k(\omega)^*$'s, which in the model is governed by the parameter λ.

A.1.2 Nested logit distributional assumptions

The results in this section are based on Chapter 4 of Train (2009), where the textbook nested logit model is introduced. Our goal is to incorporate flexibility in substitution patterns within a nest relative to across nests, while maintaining tractability. We do so with a nested logit model, which amounts to assuming the idiosyncratic shocks from the prior section $\{\varepsilon_i^1(\omega)^*, \ldots, \varepsilon_i^{\mathcal{C}}(\omega)^*, \varepsilon_i^{\mathcal{N}}(\omega)^*\}$ are distributed Generalized Extreme Value (GEV) with the following cdf,

$$\exp\left(-\sum_{n=1}^{N}\left[\sum_{k\in B_n} \exp\left(-\frac{\varepsilon_i^k(\omega)^*}{\lambda_n}\right)\right]^{\lambda_n}\right), \tag{1}$$

where N is the number of nests, B_n is the name of n-th nest, k is an option inside a nest, and $\lambda_n \in (0,1)$ is a parameter governing the degree of correlation within each nest. Lower values of λ_n indicate higher correlation, hence setting $\lambda_n = 1\ \forall n$ makes the $\varepsilon_i^k(\omega)$'s independent and turns the nested logit model into a multinomial logit model. Note the marginal distribution of each $\varepsilon_i^k(\omega)^*$ in 1 is type I EV with standard deviation $\frac{\pi}{\sqrt{6}}$ because of how we re-scaled payoffs in the preceding section. In our specific application we have N=2 and $\{B_1, B_2\} = \{\mathcal{C}, \mathcal{N}\}$. We also set $\lambda_{\mathcal{N}} = 1$ because nest $\mathcal{N}$ has a single option, rendering the correlation parameter meaningless, and we also drop the subscript for $\lambda_{\mathcal{C}}$ since we only need to estimate a single correlation parameter. Following Train (2009), under the GEV distributional assumption, the probability commodity c is chosen, conditional on the farmer choosing the agricultural nest $\mathcal{C}$, is,

$$\pi_i^{c|\mathcal{C}} = \frac{\left(r_i^c\right)^{\frac{\theta}{\lambda}}}{\sum_{c'\in\mathcal{C}}\left(r_i^{c'}\right)^{\frac{\theta}{\lambda}}}, \tag{CCP}$$

and the choice probabilities of the two nests are,

$$\pi_i^{\mathcal{C}} = \frac{\left(P_i^{\mathcal{C}}\right)^{\lambda}}{P_i^{\mathcal{N}} + \left(P_i^{\mathcal{C}}\right)^{\lambda}} \qquad \text{and} \qquad \pi_i^{\mathcal{N}} = \frac{P_i^{\mathcal{N}}}{P_i^{\mathcal{N}} + \left(P_i^{\mathcal{C}}\right)^{\lambda}}. \tag{NCP}$$

where we have defined $P_i^{\mathcal{C}} \equiv \sum_{c'\in\mathcal{C}}\left(r_i^{c'}\right)^{\frac{\theta}{\lambda}}$ as the payoff of the agricultural nest, and $P_i^{\mathcal{N}} \equiv \left(A_i^{\mathcal{N}}\right)^{\theta}$ as the payoff of the natural use nest. Technically, $\ln P_i^{\mathcal{C}}$ and $\ln P_i^{\mathcal{N}}$ are the inclusive values of each nest in the nested logit model. While θ governs the dispersion of shocks across fields (because it is technically the inverse of the parameter that drives the variance of the shocks), λ governs how correlated the shocks are across commodities within the agricultural nest. Our GEV distributional assumption implies the shocks are more correlated among agricultural commodities than between any given agricultural commodity and the natural use option. This delivers the following cross-price elasticities of land shares,

$$\frac{d\ln \pi_i^k}{d\ln p_i^{c'}} = \frac{d\ln \pi_i^k}{d\ln r_i^{c'}}\frac{\partial \ln r_i^{c'}}{\partial \ln p_i^{c'}} = \begin{cases} \left[-\theta\pi_i^{c'} + \theta\frac{(\lambda-1)}{\lambda}\pi_i^{c'|\mathcal{C}}\right]\frac{1}{\gamma_L^{c'}} & \text{for } k = c \neq c' \in \mathcal{C} \\ -\theta\pi_i^{c'}\frac{1}{\gamma_L^{c'}} & \text{for } k = \mathcal{N}. \end{cases}$$

The above implies substitution patterns are stronger within-nest than across nest, since $\left|\frac{d\ln \pi_i^c}{d\ln p_i^{c'}}\right| > \left|\frac{d\ln \pi_i^{\mathcal{N}}}{d\ln p_i^{c'}}\right|$ for $c \neq c'$. Notice the proportional substitution property holds within-nest $\mathcal{C}$ but not across nests: $\frac{d\ln \pi_i^{c|\mathcal{C}}}{d\ln p_i^{c'}} = -\frac{\theta}{\lambda}\pi_i^{c'|\mathcal{C}}\frac{1}{\gamma_L^{c'}}\ \forall c$.

A.1.3 Interpretation of within- and across-nest elasticities of substitution

First, consider the odds ratio between two commodities within the agricultural nest,

$$\ln\left(\frac{\pi_i^{c|\mathcal{C}}}{\pi_i^{c'|\mathcal{C}}}\right) = \ln\left(\frac{\pi_i^c}{\pi_i^{c'}}\right) = \frac{\theta}{\lambda}\ln\left(\frac{r_i^c}{r_i^{c'}}\right). \tag{2}$$

From the above, the parameter ratio $\frac{\theta}{\lambda}$ is the elasticity of substitution between commodities: it tells us how the relative land share between two commodities changes when their relative returns change. The reason we do not interpret $\frac{\theta}{\lambda}$ as a deforestation elasticity is because changes in the land shares π_i^c and $\pi_i^{c'}$ do not necessarily imply deforestation: these shares could be rising by stealing land shares from other commodities than from natural land. Intuitively, the substitution elasticity $\frac{\theta}{\lambda}$ is large when dispersion of shocks across fields goes to zero ($\theta \to \infty$) because marginal fields in a county become identical to the infra-marginal fields, and as a result, the supply curve becomes flat. This elasticity can also be large when the shocks are perfectly correlated across commodities ($\lambda \to 0$). With perfect correlation, the idiosyncratic shocks move in tandem across different commodities—the variance of their differences falls, implying county heterogeneity falls.

Second, consider the odds ratio between the two nests,

$$\ln\left(\frac{\pi_i^{\mathcal{C}}}{\pi_i^{\mathcal{N}}}\right) = \lambda\ln\left(\frac{P_i^{\mathcal{C}}}{(A_i^{\mathcal{N}})^{\frac{\theta}{\lambda}}}\right) \quad \text{where } P_i^{\mathcal{C}} = \sum_{c\in\mathcal{C}}\left(r_i^c\right)^{\frac{\theta}{\lambda}}, \tag{3}$$

where $P_i^{\mathcal{C}}$ is interpreted as the payoff to nest $\mathcal{C}$. The reason for defining the nest's payoff in this way is because the conventional definition for a nest's payoff in a nested logit model is the inclusive value term (Train, 2009), which in our setting takes the form of $\ln\left(\sum_{c\in\mathcal{C}}\left(r_i^c\right)^{\frac{\theta}{\lambda}}\right)$, i.e., $\ln P_i^{\mathcal{C}}$. From the above, λ indicates how agricultural land expands relative to natural land when the payoff of agriculture relative to nature increases. For this reason, λ is interpreted as a deforestation elasticity. Technically, λ is the elasticity of substitution across nests, i.e., between nature and agriculture.

A.1.4 Analytical expression of carbon tax impact

The impact of a local carbon tax on commodity c for overall land use is obtained by deriving the elasticity of the nest $\mathcal{C}$ share with respect to an individual commodity c price, since a carbon tax would be an unit-specific tax, which is akin to a price change,

$$\frac{d\ln L_i^{\mathcal{C}}}{d\ln p_i^c} = \frac{d\ln \pi_i^{\mathcal{C}}}{d\ln p_i^c} = \frac{d\ln \pi_i^{\mathcal{C}}}{d\ln r_i^c}\frac{\partial \ln r_i^c}{\partial \ln p_i^c} = \theta \pi_i^{c|\mathcal{C}}(1-\pi_i^{\mathcal{C}})\frac{1}{\gamma_L^c}.$$

We need to put a negative sign in front of the expression above in order to obtain the impact of a tax (a price decrease for the producer). The drop in agricultural land is stronger when i) θ is high, ii) the share of the individual commodity within nest $\mathcal{C}$ is especially high, and iii) the land-intensity of commodity c is small (because production depends on the homogenous non-land factors).

A.2 County-level aggregation

A.2.1 Within-nest expected returns

Only fields with high draws of $r_i^c(\omega)$ are chosen by landowners to produce c. In what follows, it is useful to define the realized payoff in field ω as the maximum payoff among all possible commodity choices within nest $\mathcal{C}$,

$$\overline{R}_i(\omega) \equiv \max_{c'\in\mathcal{C}} r_i^{c'}(\omega).$$

$\overline{R}_i(\omega)$ is a random variable, and its expected value is the expected payoff from choosing nest $\mathcal{C}$. This object has a closed form: it is the inclusive value of the nest multiplied by the nest's correlation parameter λ (Train (2009), Chapter 4). To obtain the payoff in dollar units one must then divide by the coefficient on price (in our case, θ). Applying the appropriate transformation from the type I EV to the type II EV setting gives us,

$$\mathbb{E}\left[\overline{R}_i(\omega)\right] = \left[\sum_{c'\in\mathcal{C}} (r_i^{c'})^{\frac{\theta}{\lambda}}\right]^{\frac{\lambda}{\theta}}.$$

Next, consider the expected payoff from commodity c, conditional on it being chosen (i.e., on being the highest payoff commodity), and let $f_{\overline{R}_i}(\bar{r})$ denote the probability density of $\overline{R}_i(\omega)$,

$$\begin{aligned}\int_{\bar{r}} \mathbb{E}\left[r_i^c(\omega)\middle| r_i^c(\omega)=\bar{r}\right] f_{\overline{R}_i}(\bar{r}) &= \int_{\bar{r}} \bar{r} f_{\overline{R}_i}(\bar{r}) = \mathbb{E}\left[\overline{R}_i(\omega)\right], \\ &= \left[\sum_{c'\in\mathcal{C}} (r_i^{c'})^{\frac{\theta}{\lambda}}\right]^{\frac{\lambda}{\theta}} = \frac{r_i^c}{r_i^c}\left[\sum_{c'\in\mathcal{C}} (r_i^{c'})^{\frac{\theta}{\lambda}}\right]^{\frac{\lambda}{\theta}}, \\ &= r_i^c\left[\frac{\sum_{c'\in\mathcal{C}} (r_i^{c'})^{\frac{\theta}{\lambda}}}{(r_i^c)^{\frac{\theta}{\lambda}}}\right]^{\frac{\lambda}{\theta}} = r_i^c\left(\pi_i^{c|\mathcal{C}}\right)^{-\frac{\lambda}{\theta}} > r_i^c = \mathbb{E}\left[r_i^c(\omega)\right].\end{aligned}$$

The last line above is the selection effect: the average return of fields allocated to produce commodity c is higher than the unconditional mean return in c. As $\pi_i^{c|\mathcal{C}}$ grows the average return falls because the fields that are less productive in commodity c begin to be drawn into production. Also, as dispersion across fields goes to zero ($\theta \to \infty$), all fields become identical and the selection effect disappears: the conditional expected return converges to the unconditional one.

A.2.2 Across-nest expected productivity

The relevant expectation over productivity needs to take into account the payoff from the natural option as well, which the previous section ignores. Therefore, we now consider the maximum payoff across all nests,

$$\overline{\overline{R}}_i(\omega) \equiv \max_{\{1, \dots\, c\,, \dots\,, C\,, \mathcal{N}\}} \left[r^1_{it}(\omega)\,,\, \dots\,,\, r^C_{it}(\omega)\,,\, A^{\mathcal{N}}_{it}(\omega) \right].$$

Farrokhi and Pellegrina (2023) show how to derive the expected productivity of a nested discrete choice problem with the same structure as mine: an upper nest where farmers choose crops, and a lower nest where they choose a technology type conditional on the crop choice. They have a GEV distribution for the idiosyncratic shocks with two parameters: θ_1 controls the dispersion of land productivity across crops, and θ_2 controls the dispersion of productivity across technologies within a crop. The ratio θ_2/θ_1 governs the correlation of the shocks across technologies within a crop (i.e., the correlation within the lower nest). My model directly maps to their framework by specifying their upper nest as my model's choice between $\mathcal{N}$ and $\mathcal{C}$, and their lower nest as my model's choice of commodity within nest $\mathcal{C}$ (if nest $\mathcal{N}$ is chosen there is no second choice since there is only one choice in nest $\mathcal{N}$). Their GEV parameters can be mapped to mine as $\theta_2 = \theta/\lambda$ and $\theta_1 = \theta$. Hence, $\frac{\theta}{\lambda}$ governs the dispersion of shocks across commodities within nest $\mathcal{C}$, while θ governs the dispersion of shocks across nests. This parameter mapping also implies that my correlation parameter λ equals the ratio of their two dispersion parameters, i.e., $\lambda = \theta_1/\theta_2$. Hence, as θ_1 falls relative to θ_2, and thus λ falls, correlation of productivity draws in the lower nest increases (in my case, the lower nest is nest $\mathcal{C}$, while in their case it is the stage of technology choices within a crop). Thus, by adjusting the notation from Farrokhi and Pellegrina (2023) to my framework, the expected productivity of commodity c, conditional on it being chosen (i.e., on being the highest payoff across all nests) is,

$$\begin{aligned}
\mathbb{E}\left[r^c_i(\omega) \middle| r^c_i(\omega) = \overline{\overline{R}}_i(\omega) \right] &= r^c_i \left(\pi^{c|\mathcal{C}}_i \right)^{-\frac{\lambda}{\theta}} \left(\pi^{\mathcal{C}}_i \right)^{-\frac{1}{\theta}} \\
&> r^c_i \left(\pi^{c|\mathcal{C}}_i \right)^{-\frac{\lambda}{\theta}} = \mathbb{E}\left[r^c_i(\omega) \middle| r^c_i(\omega) = \overline{R}_i(\omega) \right] \\
&> r^c_i = \mathbb{E}\left[r^c_i(\omega) \right].
\end{aligned}$$

In the second to last line, notice the expected productivity is now decreasing in both the within-nest share of commodity c and the nest $\mathcal{C}$ share. As more fields are allocated to commodity c specifically and nest $\mathcal{C}$ broadly, these fields are the ones with both i) lower yields, and ii) lower commodity-relative-to-nature payoff. This expectation is higher than the within-nest $\mathcal{C}$ expectation because we are now taking into account a second source of selection in the upper nest.

A.2.3 County-level quantities

To obtain quantities at the county level we directly follow the approach in Farrokhi and Pellegrina (2023). First, note the land share of revenue is equal to the land rents,

$$\gamma^c_L p^c_{it} Q^c_{it}(\omega) = r^c_{it}(\omega) \implies Q^c_{it}(\omega) = \frac{r^c_{it}(\omega)}{p^c_{it}\gamma^c_L}.$$

Therefore, the county-level expected output is given by,

$$\begin{aligned}\mathbb{E}\left[Q_{it}^{c}(\omega)\middle| r_{it}^{c}(\omega)=\overline{\overline{R}}_{it}(\omega)\right] &= \frac{1}{p_{it}^{c}\gamma_{L}^{c}}\times\mathbb{E}\left[r_{it}^{c}(\omega)\middle| r_{it}^{c}(\omega)=\overline{\overline{R}}_{it}(\omega)\right] \\ &= \frac{r_{it}^{c}}{p_{it}^{c}\gamma_{L}^{c}}\left(\pi_{it}^{c|\mathcal{C}}\right)^{-\frac{\lambda}{\theta}}\left(\pi_{it}^{\mathcal{C}}\right)^{-\frac{1}{\theta}}.\end{aligned}$$

The above is the expected output on an infinitesimally small field ω. To obtain county-level output we scale up by the amount of land allocated to commodity c,

$$\begin{aligned}Q_{it}^{c} &= \mathbb{E}\left[Q_{it}^{c}(\omega)\middle| r_{it}^{c}(\omega)=\overline{\overline{R}}_{it}(\omega)\right]\times\pi_{it}^{c}\bar{L}_{i} \\ &= \frac{r_{it}^{c}}{p_{it}^{c}\gamma_{L}^{c}}\left(\pi_{it}^{c|\mathcal{C}}\right)^{-\frac{\lambda}{\theta}}\left(\pi_{it}^{\mathcal{C}}\right)^{-\frac{1}{\theta}}\times\pi_{it}^{c}\bar{L}_{i} \\ &= \frac{r_{it}^{c}}{p_{it}^{c}\gamma_{L}^{c}}\left(\pi_{it}^{c|\mathcal{C}}\right)^{1-\frac{\lambda}{\theta}}\left(\pi_{it}^{\mathcal{C}}\right)^{1-\frac{1}{\theta}}\bar{L}_{i}.\end{aligned}$$

A.3 Consumer CES demand problem

A destination j consumer maximizes her utility U_{jt} by choosing C_{ijt}^{k}: how much to consume of product k from county i given prices p_{ijt}^{k} and a total budget of X_{jt}. Here, X_{jt} is the destination j income spent on agricultural goods, which I assume is exogenous since this is a single-industry model. Formally, the consumer j problem is,

$$\max_{C_{ijt}^{c}} \quad U_{jt} \quad s.t. \quad \sum_{c}\sum_{i} p_{ijt}^{c}C_{ijt}^{c}\leq X_{jt}, \quad \text{where} \quad U_{jt}\equiv\left(\sum_{k}(a_{jt}^{k})^{\frac{1}{\eta_u}}(C_{jt}^{k})^{\frac{\eta_u-1}{\eta_u}}\right)^{\frac{\eta_u}{\eta_u-1}},$$

where $C_{jt}^{k}=\left(\sum_{n}(a_{njt}^{k})^{\frac{1}{\eta_m}}(C_{njt}^{k})^{\frac{\eta_m-1}{\eta_m}}\right)^{\frac{\eta_m}{\eta_m-1}}$ is consumption of good k aggregated across source nations indexed n and $C_{njt}^{k}=\left(\sum_{i'\in n}(a_{i'jt}^{k})^{\frac{1}{\eta_l}}(C_{i'jt}^{k})^{\frac{\eta_l-1}{\eta_l}}\right)^{\frac{\eta_l}{\eta_l-1}}$ is consumption of good k aggregated across source counties indexed i' belonging to nation n. The η_l, η_m, and η_u terms are the lower, middle, and upper level elasticities of substitution and the a terms are preference shifters at each level. Moreover, it will be useful to define price indices at each level of the demand system as $P_{njt}^{k}\equiv\left(\sum_{i'\in n}a_{i'jt}^{k}(p_{i'jt}^{k})^{1-\eta_l}\right)^{\frac{1}{1-\eta_l}}$, $P_{jt}^{k}\equiv\left(\sum_{n}a_{njt}^{k}(P_{njt}^{k})^{1-\eta_m}\right)^{\frac{1}{1-\eta_m}}$ and $P_{jt}\equiv\left(\sum_{c}a_{jt}^{k}(P_{jt}^{k})^{1-\eta_u}\right)^{\frac{1}{1-\eta_u}}$.

Lower level. Combining the first order conditions of the consumer problem for two origin-specific products C_{ijt}^{k} and $C_{i'jt}^{k}$ with $i,i'\in n$, and then summing for all $i'\in n$ we obtain,

$$C_{ijt}^{k}=a_{ijt}^{k}\left(\frac{p_{ijt}^{k}}{P_{njt}^{k}}\right)^{-\eta_l}\frac{X_{njt}^{k}}{P_{njt}^{k}}, \tag{4}$$

where $X^k_{njt} \equiv \sum_{i' \in n} p^k_{i'jt} C^k_{i'jt}$ is total expenditure on product k from nation n. To obtain the lower level estimating equation, multiply both sides of 4 by p^k_{ijt}, bring the X^k_{njt} term to the LHS, take logs,

$$\ln\left(\frac{X^k_{ijt}}{X^k_{njt}}\right) = (1-\eta_l)\ln\left(p^k_{ijt}\right) - \ln\left(\sum_{i' \in n} a^k_{i'jt}(p^k_{i'jt})^{1-\eta_l}\right) + \ln\left(a^k_{ijt}\right) \quad \forall i \in n. \tag{5}$$

Middle level. Plug in 4 into the consumer's FOC for C^k_{ijt}, evaluate it for n and n', and then sum across all n'. This delivers the following expression,

$$\frac{X^k_{njt}}{X^k_{jt}} = \frac{a^k_{njt}\left(P^k_{njt}\right)^{1-\eta_m}}{\sum_{n'} a^k_{n'jt}\left(P^k_{n'jt}\right)^{1-\eta_m}}$$

Taking logs of the above delivers our estimating equation for the middle level,

$$\ln\left(\frac{X^k_{njt}}{X^k_{jt}}\right) = (1-\eta_m)\ln\left(P^k_{njt}\right) - \ln\left(\sum_{n'} a^k_{n'jt}(P^k_{n'jt})^{1-\eta_m}\right) + \ln\left(a^k_{njt}\right). \tag{6}$$

Upper level. Plug in 4 and C^k_{njt}, into the consumer's FOC for C^k_{ijt}, evaluate it for k and k', and sum across all k'. This delivers the following expression,

$$\frac{X^k_{jt}}{X_{jt}} = \frac{a^k_{jt}\left(P^k_{jt}\right)^{1-\eta_u}}{\sum_{k'} a^{k'}_{jt}\left(P^{k'}_{jt}\right)^{1-\eta_u}}.$$

Taking logs of the above delivers our estimating equation for the upper level,

$$\ln\left(\frac{X^k_{jt}}{X_{jt}}\right) = (1-\eta_u)\ln\left(P^k_{jt}\right) - \ln\left(\sum_{k'} a^{k'}_{jt}\left(P^{k'}_{jt}\right)^{1-\eta_u}\right) + \ln\left(a^k_{jt}\right). \tag{7}$$

A.3.1 Price-index level expression for demand

Notice we can write C^c_{ijt} from 4 in terms of price indices at each level and the total expenditure X_{jt},

$$C^k_{ijt} = a^k_{ijt}\left(\frac{p^k_{ijt}}{P^k_{njt}}\right)^{-\eta_l} a^k_{njt}\left(\frac{P^k_{njt}}{P^k_{jt}}\right)^{-\eta_m} a^k_{jt}\left(\frac{P^k_{jt}}{P_{jt}}\right)^{-\eta_u} \frac{X_{jt}}{P_{jt}} \quad \forall i \in n, \tag{8}$$

which is the expression in the main text.

A.4 Extension with firm heterogeneity and entry/exit

Entry with homogenous firms. Consider an extension of the intermediary model that allows for entry, but maintains the firm homogeneity assumption. We assume that in each market i there is a fixed entry cost ϕ_i^c. Potential entrants enter until the incumbent profits equal the entry cost,

$$\bar{p}_{it}^k f(q_{it}^c;k) - \underbrace{\bar{p}_{it}^k f'(q_{it}^c;k)\mu_{it}^c(q_{it}^c, N_{it}^c)}_{=p_{it}^c} q_{it}^c = \phi_i^c.$$

To simplify, the expression above, consider the case of linear technology $f(q_{it}^c;k) = \alpha_c q_{it}^c$,

$$\bar{p}_{it}^k \alpha_c q_{it}^c \left[1 - \mu_{it}^c(q_{it}^c, N_{it}^c)\right] = \phi_i^c.$$

The key empirical challenge is obtaining estimates of the entry cost ϕ_i^c. One approach is to use the number of firms at one point in time to back out values for ϕ_i^c, and in counterfactuals assume ϕ_i^c is fixed. However, these assumptions are quite strong and not too different from directly fixing the number of firms, since ϕ_i^c is backed out from that number at some point in time and kept fixed. It is not obvious this approach is more credible than the baseline assumption of a fixed number of firms, especially since there isn't much observed entry/exit in the data to begin with. For this reason, I use the baseline model from the main text, which is upfront about the difficulty of dealing with entry. Apart from introducing empirical issues, it seems unlikely that adding exit/entry would change the paper's qualitative insights in terms of the transmission of environmental policy. Recall that the qualitative insights from introducing market power in the baseline model rely solely on incomplete pass-through. To the extent that adding exit/entry continues to deliver incomplete pass-through, the qualitative results of the paper would therefore remain the same. Given our main goal is not to measure the welfare impact of market power per se, but simply how it interferes with the transmission of environmental policy, allowing for exit/entry is unlikely to be crucial for the paper's qualitative results.

No entry, but heterogeneous firms. With heterogeneous firms, we would have to introduce an additional elasticity of substitution across firms in the farmers' choice problem. To credibly estimate this elasticity, we would need plausibly exogenous demand shifters at the firm level as instruments, which are hard to come by. As already mentioned, it is important to stress that the qualitative insights from introducing market power in the baseline model rely solely on incomplete pass-through. To the extent that adding firm heterogeneity continues to deliver incomplete pass-through, the qualitative results of the paper would remain the same. Thus, given our main goal is not to measure the welfare impact of market power per se, but simply how it interferes with the transmission of environmental policy, allowing for firm heterogeneity is unlikely to be crucial for the paper's qualitative results.

Entry with heterogeneous firms. If one combines firm heterogeneity with entry then we can have selection of highly productive firms into the market, and thus obtain the typical extensive-margin efficiency gains emphasized by the long-standing firm heterogeneity literature. In general, more productive firms can be more environmentally friendly if they use resources more efficiently. While including these ingredients would allow us to have productivity effects on agribusiness firms, it would seem unlikely this is first order for environmental outcomes in our agricultural setting since it is the farmers, not the agribusiness firms, who ultimately make the environmentally relevant decisions. For other industries though, such as manufacturing, the extensive-margin efficiency gains could have important environmental consequences if the more productive firms use fossil fuels more efficiently, for example. In any case, and as mentioned in the prior paragraph, to the extent that adding firm heterogeneity and exit continues to deliver incomplete pass-through, the qualitative results of the paper are likely to remain the same.

A.5 Decentralizing the first-best allocation

Let variables corresponding to the first-best allocations be denoted by superscript *. For the equilibrium outcomes in the imperfectly competitive case we remove superscript *. First, we set the forest subsidy at the same Pigouvian rate as in the first best $s_i^{\mathcal{N}} = s_i^{\mathcal{N}*}$, so that landowners fully internalize the land use change component of emissions.

Second, we need farmers to internalize that commodities vary in their non-LUC emissions. Hence, we want to choose a commodity-specific output tax t_i^c that decentralizes the first-best. That is, we need farmers to receive the same after-tax price as in the first-best so that they choose $L_i^c = L_i^{c*}$. The following condition guarantees this,

$$p_i^c - t_i^c = p_i^{c^*} - t_i^{c*}, \tag{9}$$

where p_i^c is the pre-tax equilibrium farm-gate price in the imperfectly competitive world. We also need consumers to demand the same amount as in the first-best, otherwise we don't have an equilibrium. The following condition guarantees this,

$$p_{ij}^k = p_{ij}^{k*} \implies \frac{p_i^c}{\mu_i^c}\frac{\tau_{ij}}{m_i^{ck}} = p_i^{c*}\frac{\tau_{ij}}{m_i^{ck}} \implies p_i^c = p_i^{c*}\mu_i^c(L_i^{c*}). \tag{10}$$

In the last equality, notice that given L_i^{c*} we can compute μ_i^c since it is only a function of land shares. Therefore, once we have p_i^{c*} and L_i^{c*}, we can pin down the equilibrium farm-gate (pre-tax) price p_i^c in the imperfectly competitive world that would deliver the same first-best allocation. Finally, we can plug condition 10 in 9 to pin down the required output tax for the farmers to produce the same as in the first-best,

$$t_i^c = t_i^{c*} - p_i^{c^*}[1 - \underbrace{\mu_i^c(L_i^{c*})}_{<1}].$$

Notice that $t_i^c < t_i^{c*}$: imperfect competition is already bringing quantities part of the way down to the first-best, so the required output tax is less than in the perfectly competitive case. Also, it could even be that t_i^c is negative, so that we would need to subsidize output. This would only occur if the environmental externality is small relative to the market power distortion, so that the imperfectly competitive equilibrium delivers quantities that are below the social optimum even after considering environmental costs.

Appendix B. Data

B.1 Cattle productivity

FAO-GAEZ reports yields in units of tonnes/ha for crops, but reports a pasture productivity index rather than a physical yield for cattle. Hence, we need to transform the index into cattle "yields per ha". Let x_i denote the FAO-GAEZ pasture productivity index in location i. Let h_i denote the location's stocking rate (number of cattle head per ha, obtained from agricultural censuses). Let $\hat{h}(x_i)$ denote an estimate of the stocking rate that can be sustained in a location with pasture index x_i. I obtain this estimate by projecting x_i onto h_i, with results shown in Table 1.

Table 1: Correlation between stocking rates and pasture productivity index.

	Dependent variable: Cattle head/ha			
	(1)	(2)	(3)	(4)
intercept	0.570***	0.660***	0.656***	0.599***
	(0.049)	(0.062)	(0.085)	(0.082)
pasture yield (index)	0.766***	0.888***	0.809***	0.885***
	(0.067)	(0.064)	(0.062)	(0.061)
maize yield (t/ha)		-0.033***		-0.044***
		(0.009)		(0.011)
soybean yield (t/ha)			-0.051*	0.055
			(0.028)	(0.033)
Num.Obs.	5,265	5,265	5,265	5,265
R2	0.150	0.168	0.154	0.170

Notes: All specifications are weighted by pasture area (ha). HC-robust SE in parenthesis.

When constructing the h_i variable I restrict the sample to calf data, since the goal is to measure spatial variation in pasture's stocking potential for breeding, and not to pick up variation in intermediate stages such as backgrounding or finishing. That is, we want to link $\hat{h}(x_i)$ to the productivity of breeders (that rely almost exclusively on pasture) and not the finishing establishments (that rely more on feed and other intermediate inputs). Finally, I use data on calf weights w_i (tonnes/head) from CEPEA to go from head/ha to tonnes/ha. The end result is a measure of the observable component of cattle productivity (in tonnes/ha), constructed as $z_i^b = \hat{h}(x_i)w_i$.

B.2 Domestic trade data

To map domestic trade flows I use supply chain data from TRASE. This data links production in domestic counties to exporting firms, logistics hubs/ports, and importing firms in foreign markets, for the major agricultural commodities of Argentina and Brazil (beef, soybeans, maize). For Brazil, soybean is available for 2004-2020, beef is available for 2010-2020, maize is available for 2015-2017. For Argentina, data is available for soybean, from 2015-2019. While we are missing agribusiness data for commodities such as wheat or rice, the three commodities we do have agribusiness data for account for roughly 85% of all agricultural land in both countries (with pasture land alone accounting for 70%).

Methodology. In order to construct the supply chain data, TRASE uses an enhanced form of material flow analysis: the "Spatially Explicit Information on Production to Consumption Systems" methodology (SEI-PCS), described in Godar, Persson, Tizado and Meyfroidt (2015). The primary inputs for the SEI-PCS methodology are customs documents (e.g., customs declarations, bills of lading) that declare the volume and FOB value of a commodity shipment, the identity of the exporting agribusiness firm, the port of export, the identity of the importing firm, and the destination country. The way that SEI-PCS links the shipments from the port of export back to the source county is by combining the customs documents with the agribusiness firm's domestic asset and tax information (e.g, information of slaughterhouses or soybean crushing facilities owned by the exporting firm). Further details on the methodology can be found at trase.earth/methodology. To the best of my knowledge, this is the most comprehensive source of domestic supply chain data available, despite not being perfect. An example of recent published work using the data is Rajão, Soares-Filho, Nunes, Börner, Machado, Assis, Oliveira, Pinto, Ribeiro, Rausch et al. (2020).

Key variables. Trade flows between source county and destination country are reported in "raw commodity-equivalent" quantities (weight) and value (port-of-export FOB). The quantities are reported in this "commodity-equivalent" way because the commodities are not always exported in the same format. For example, in the case of beef the data specifies whether the exports are "raw" live animals, de-boned frozen beef, offals, etc. In the case of maize or soybeans, the data specifies whether the exports are in raw grain form or partially processed into meal. Reported quantities therefore reflect the equivalent raw-commodity weight attributable to the exported format, based on a linear conversion factor. Finally, the data also specifies the location of the logistics hub (e.g., ports and inland consignment centers) through which the commodity is shipped.

Validation to conventional agricultural trade datasets. To get a sense of the quality and representativeness of the TRASE data, I compare it to conventional aggregated data sources, such as FAOSTAT. Figure 1 shows how the national-level exports from Argentina and Brazil implied by the dis-aggregated TRASE data line up with FAOSTAT. The left panel of the figure shows how exports correlate across data sources (i.e., each point is a producing nation-commodity-year observation). The right panel shows how bilateral export pairs between each producing nation (Argentina or Brazil) and destination correlate across data sources (i.e., each point is a producing nation-destination-commodity-year observation). For both total export levels as well as the bilateral pairs, TRASE aggregates are highly correlated with FAOSTAT, with an R-squared above 99% and a slope of best fit near 1.

Figure 1: Validation of TRASE micro-data with FAOSTAT aggregate data.

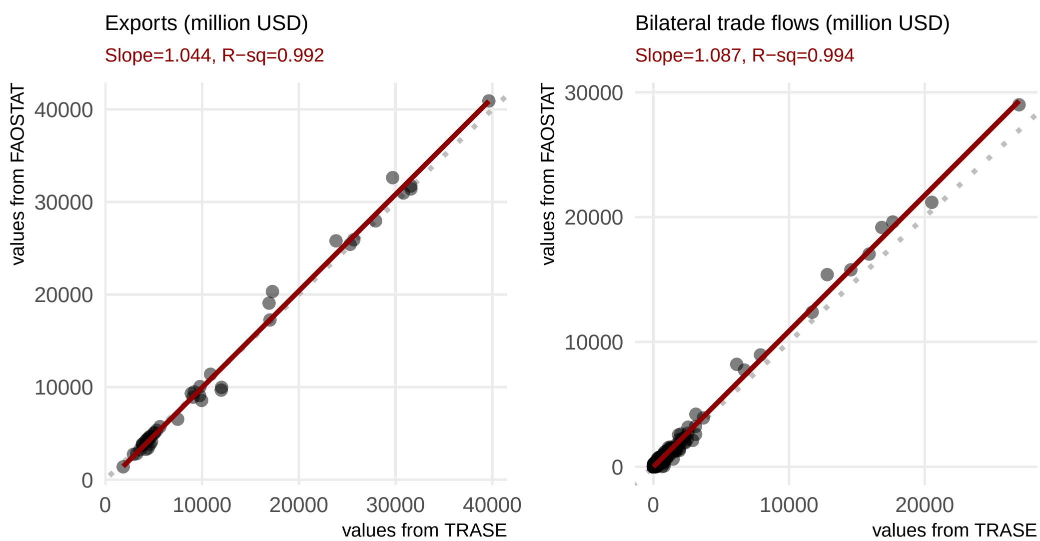

Notes: The left panel shows how export values for each producing nation (Argentina or Brazil) correlate across the two data sources (i.e., each point is a producing nation-commodity-year observation). The right panel shows how bilateral export values between each producing nation (Argentina or Brazil) and a destination nation correlate across the two data sources (ie., each point is a producing nation-destination nation-commodity-year observation).

B.2.1 Persistence analysis

Persistence of agribusiness firm presence over time. The presence and market share of an individual firm in an upstream market is persistent during the sample period. To show this pattern I estimate the following AR(1) regression,

$$y^c_{fit} = \rho y^c_{fit-1} + \varepsilon^c_{fit},$$

where y^c_{fit} is either i) the market share in commodity c of firm f in county i at time t, or ii) a dummy variable indicating whether firm f is actively sourcing commodity c from county i at time t. The results for both outcomes are displayed in Tables 2-3, respectively, and suggest high persistence in specific upstream markets, regardless of how granular the upstream markets are defined.

Table 2: Persistence of a firm's market share in an upstream market.

	AR(1)	AR(1)	AR(1)	AR(1)
ρ	0.533***	0.618***	0.659***	0.695***
	(0.004)	(0.008)	(0.013)	(0.031)
Upstream market definition	AMC/County	Mesoregion	Microregion	State/Province
Observations	33,790,006	5,523,836	1,690,518	379,080
R2	0.283	0.383	0.432	0.488

Notes: results are from an AR(1) model, $y^c_{fit} = \rho y^c_{fit-1} + \varepsilon^c_{fit}$, where y^c_{fit} is the market share in commodity c of firm f in upstream market i at time t. The data covers Argentina and Brazil, and all agribusiness firms operating in beef, soybean, and maize markets. Each column defines upstream markets at different levels of spatial aggregation. HC-robust SE in parenthesis.

Table 3: Persistence of a firm's presence in an upstream market.

	(1)	(2)	(3)	(4)
ρ	0.660***	0.660***	0.666***	0.681***
	(0.002)	(0.003)	(0.004)	(0.008)
Upstream market definition	AMC/County	Mesoregion	Microregion	State/Province
Observations	33,790,006	5,523,836	1,690,518	379,080
R2	0.400	0.398	0.413	0.417

Notes: results are from an AR(1) model, $y^c_{fit} = \rho y^c_{fit-1} + \varepsilon^c_{fit}$, where y^c_{fit} is a dummy indicator that takes a value of 1 if firm f is sourcing commodity c from upstream market i at time t. The data covers Argentina and Brazil, and all agribusiness firms operating in beef, soybean, and maize markets. Each column defines upstream markets at different levels of spatial aggregation. HC-robust SE in parenthesis.

Persistence of an upstream location's exposure to foreign consumer markets. We can repeat the same analysis of persistence, but for foreign consumer markets instead of agribusiness firms. To show these patterns I estimate the following AR(1) regression,

$$y^c_{jit} = \rho y^c_{jit-1} + \varepsilon^c_{jit},$$

Table 4: Persistence of a destination's market share in an upstream market.

	(1)	(2)	(3)	(4)
ρ	0.648***	0.763***	0.824***	0.811***
	(0.004)	(0.007)	(0.010)	(0.031)
Upstream market definition	AMC/County	Mesoregion	Microregion	State/Province
Observations	373,674	64,568	22,890	4,060
R2	0.407	0.570	0.674	0.641

Notes: results are from an AR(1) model, $y^c_{jit} = \rho y^c_{jit-1} + \varepsilon^c_{jit}$, where y^c_{jit} is the market share in commodity c of destination j in upstream market i at time t. The data covers Argentina and Brazil, and beef, soybean, and maize markets. Each column defines upstream markets at different levels of spatial aggregation. HC-robust SE in parenthesis.

Table 5: Persistence of a destination's presence in an upstream market.

	(1)	(2)	(3)	(4)
ρ	0.689***	0.694***	0.678***	0.743***
	(0.002)	(0.005)	(0.008)	(0.018)
Upstream market definition	AMC/County	Mesoregion	Microregion	State/Province
Observations	373,674	64,568	22,890	4,060
R2	0.472	0.480	0.459	0.553

Notes: results are from an AR(1) model, $y^c_{jit} = \rho y^c_{jit-1} + \varepsilon^c_{jit}$, where y^c_{jit} is a dummy indicator that takes a value of 1 if destination j is sourcing commodity c from upstream market i at time t. The data covers Argentina and Brazil, and beef, soybean, and maize markets. Each column defines upstream markets at different levels of spatial aggregation. HC-robust SE in parenthesis.

where y^c_{jit} is either i) the market share in commodity c of destination j in county i at time t, or ii) a dummy variable indicating whether destination j is actively sourcing commodity c from county i at time t. The results for both outcomes are displayed in Tables 4-5, indicating high persistence in specific upstream markets, regardless of how granular the upstream markets are defined.

Appendix C. Estimation

C.1 Role of the nesting structure

Estimating a non-nested multinomial model amounts to imposing $\lambda = 1$ and estimating a single parameter θ from variation between commodities and natural use simultaneously. The multinomial model mixes variation within and across nests to deliver a single elasticity: the land-use change elasticity and the substitution elasticity between commodities are one and the same,

$$\frac{d \ln \pi^c_{it}}{d \ln p^c_{it}} = \theta(1 - \pi^c_{it})/\gamma^c_L.$$

The nested model's objective is to allow changes in a commodity's acreage to be decomposed into both margins:

$$\frac{d\ln \pi_{it}^{c}}{d\ln p_{it}^{c}} = \underbrace{\frac{d\ln \pi_{it}^{c|\mathcal{C}}}{d\ln p_{it}^{c}}}_{\text{commodity substitution: } \frac{\theta}{\lambda}(1-\pi_{it}^{c|\mathcal{C}})/\gamma_L^c} + \underbrace{\frac{d\ln \pi_{it}^{\mathcal{C}}}{d\ln p_{it}^{c}}}_{\text{land conversion: } \frac{\theta}{\lambda}\pi_{it}^{c|\mathcal{C}}\times\lambda(1-\pi_{it}^{\mathcal{C}})/\gamma_L^c} .$$

The first term is identical to the land-use change elasticity implied by Ricardian agricultural trade models and indicates how crop c acreage increases by stealing land shares from other crops. The second term tells us how crop c land use increases by stealing land shares from natural land. Total agricultural land responds to overall agricultural returns, measured as $P_{it}^{\mathcal{C}}$, as follows:

$$\frac{d\ln \pi_{it}^{\mathcal{C}}}{d\ln P_{it}^{\mathcal{C}}} = \lambda(1-\pi_{it}^{\mathcal{C}}),$$

By separating the two margins, the nested model adds flexibility so that substitution elasticities across commodities are estimated from variation within the agricultural nest while land-use change elasticities are estimated from variation across nests.

C.2 Construction of inclusive value for across-nest estimation

To estimate the across-nest equation 8 of the main text we need to construct the inclusive value term $\ln P_{it}^{\mathcal{C}}$, where $P_{it}^{\mathcal{C}} \equiv \sum_{c\in\mathcal{C}} \left(r_{it}^{c}\right)^{\frac{\theta}{\lambda}}$. Hence, we need to know the r_{it}^{c} terms, which consist of observables (agronomic yields z_{it}^{c}, farm-gate prices p_{it}^{c}, and factor shares γ_L^c) as well unobservables,

$$r_{it}^{c} = z_{it}^{c}\,(p_{it}^{c})^{\frac{1}{\gamma_L^c}}\,\tilde{\zeta}_{it}^{c} \quad \text{where} \quad \tilde{\zeta}_{it}^{c} \equiv (w_{Hit})^{-\frac{\gamma_H^c}{\gamma_L^c}}\,(w_{Mit})^{-\frac{\gamma_M^c}{\gamma_L^c}}\,\zeta_{it}^{c}. \tag{11}$$

Above, I've labeled $\tilde{\zeta}_{it}^{c}$ as the unobservable part of r_{it}^{c}: it is a combination of the unobservable land productivity shocks ζ_{it}^{c} and unobserved non-land input prices (w_{Mit}, w_{Hit}). For estimation purposes we don't need to separate ζ_{it}^{c} from $\tilde{\zeta}_{it}^{c}$, but for the model simulations that come later we will (this is explained later on in Appendix D.1, where I describe how I simulate the model). Below, I outline the steps to estimate the across-nest equation 8 of the main text. One of these steps involves using data on observed yields to calibrate the unobserved $\tilde{\zeta}_{it}^{c}$ terms that enter the inclusive value.

1. Estimate the within-nest substitution elasticities $\frac{\theta}{\lambda}$ from equation 7 of the main text.
2. Given an estimate of $\frac{\theta}{\lambda}$, back out relative values of $\tilde{\zeta}_{it}^{c}$ from equation 7 of the main text,

$$\ln\left(\frac{\tilde{\zeta}_{it}^{c}}{\tilde{\zeta}_{it}^{c'}}\right) = \frac{\lambda}{\theta}\ln\left(\frac{\pi_{it}^{c}}{\pi_{it}^{c'}}\right) - \ln\left(\frac{\left(p_{it}^{c}\right)^{\frac{1}{\gamma_L^c}} z_{it}^{c}}{\left(p_{it}^{c'}\right)^{\frac{1}{\gamma_L^{c'}}} z_{it}^{c'}}\right) \tag{12}$$

3. To construct the inclusive value terms we need to pin down the level of $\tilde{\zeta}_{it}^{c}$ $\forall c$. Notice it suffices to pin down the level of $\tilde{\zeta}_{it}^{c}$ for one commodity, because we can then use the relative values from 12 to obtain the rest. Hence, for that one commodity c, we can use the model's

implied yields to backout the value of $\tilde{\zeta}^c_{it}$ that would rationalize an observed yield of $\frac{Q^c_{it}}{L^c_{it}}$,

$$\frac{Q^c_{it}}{L^c_{it}} = \frac{\left(p^c_{it}\right)^{\frac{1}{\gamma^c_L}-1}}{\gamma^c_L} z^c_{it}\tilde{\zeta}^c_{it}\left(\pi^{c|\mathcal{C}}_{it}\right)^{-\frac{\lambda}{\theta}}\left(\pi^{\mathcal{C}}_{it}\right)^{-\frac{1}{\theta}} \implies \tilde{\zeta}^c_{it} = \frac{\gamma^c_L}{\left(p^c_{it}\right)^{\frac{1}{\gamma^c_L}-1}}\frac{Q^c_{it}}{L^c_{it}}\frac{1}{z^c_{it}}\left(\pi^{c|\mathcal{C}}_{it}\right)^{\frac{\lambda}{\theta}}\left(\pi^{\mathcal{C}}_{it}\right)^{\frac{1}{\theta}}.$$

Thus, we have expressed $\tilde{\zeta}^c_{it}$ as a function of observables: $\frac{Q^c_{it}}{L^c_{it}}$ are the observed yields we want the model to match, $\pi^{c|\mathcal{C}}_{it}$ and $\pi^{\mathcal{C}}_{it}$ are observed land shares, $\frac{\lambda}{\theta}$ is known from the within-nest estimation step, and z^c_{it} is the agronomic FAO-GAEZ data.[1]

C.3 Markdown analysis

C.3.1 Robustness of model-implied markdowns to definition of upstream market

Recall the model-implied markdown is given by $\mu^c_i \equiv \left(1 + \frac{\partial \ln p^c_i}{\partial \ln Q^c_i}\frac{1}{N^c_i}\right)^{-1} < 1$. Markdowns are a function of the number of firms N^c_i and the supply elasticity $\frac{\partial \ln p^c_i}{\partial \ln Q^c_i}$. The price-elasticity of supply is $\frac{\partial \ln Q^c_{it}}{\partial \ln p^c_{it}} = \left[\left(\frac{\theta}{\lambda}-1\right)\left(1-\pi^{c|\mathcal{C}}_{it}\right) + (\theta-1)\pi^{c|\mathcal{C}}_{it}\left(1-\pi^{\mathcal{C}}_{it}\right)\right]\frac{1}{\gamma^c_L} + \left(\frac{1}{\gamma^c_L}-1\right)$: a function of land use substitution parameters (θ, λ), factor intensities (γ^c_L), and land use shares (π^c_{it}). Notice that the larger we define the upstream market, the more agribusiness firms N^c_i will be operating there, and markdowns μ^c_i will converge up towards 1. At the same time, as we change the upstream market definition, the supply elasticity also changes because the land use shares change. In my baseline estimation exercise I use Brazilian AMCs as my upstream market definition, since using the smallest spatial unit possible is useful for estimation precision.

Figure 2 compares the values of N^c_i and μ^c_i under my baseline to that implied by two alternative upstream market definitions: microregions and mesoregions. The two top panels show markdowns are indeed closer to 1 the larger we define the upstream market: the intercept in the scatter plots are shifted up. However, the spatial distribution of markdowns remains fairly stable, as reflected by the high R^2: the locations with the largest markdowns are the same regardless of how the upstream market is defined, even though the average level of markdowns is scaled up. Recall the qualitative mechanisms of the paper depend on the spatial correlation of markdowns rather than their average level—it is the fact that markdowns are wider in emissions-intense regions that delivers the paper's qualitative results.

The two bottom panels of Figure 2 show the number of firms operating in an upstream market is mechanically larger as we increase the size of the upstream market—the intercept in the scatter plots are shifted up. Less mechanically, notice the slope of line of best fit is greater than 1: as we increase the size of the spatial unit, it is the competitive locations (according to the baseline AMC definition) that disproportionately increase the number of firms. By making the competitive regions even more competitive as we increase the size of the spatial unit, the qualitative policy results are actually exacerbated because the correlation between market power and high emissions intensity becomes stronger. Finally, Figure 3 shows that emissions intensities are negatively correlated with model-implied markdowns—the places where farm-gate prices are marked down most are the emissions-intense ones—and this holds across definitions of upstream market.

[1] Notice we also need to know θ. Hence, we make a guess for λ, which restricts θ through our estimate of $\frac{\theta}{\lambda}$ from step 1. This allows us to pin down $\tilde{\zeta}^c_{it}$ and estimate 8 of the main text, from which we obtain an estimate for λ. We therefore iterate until our guess of λ converges to the estimated λ.

Figure 2: Comparison of number of firms and markdowns across upstream market definitions.

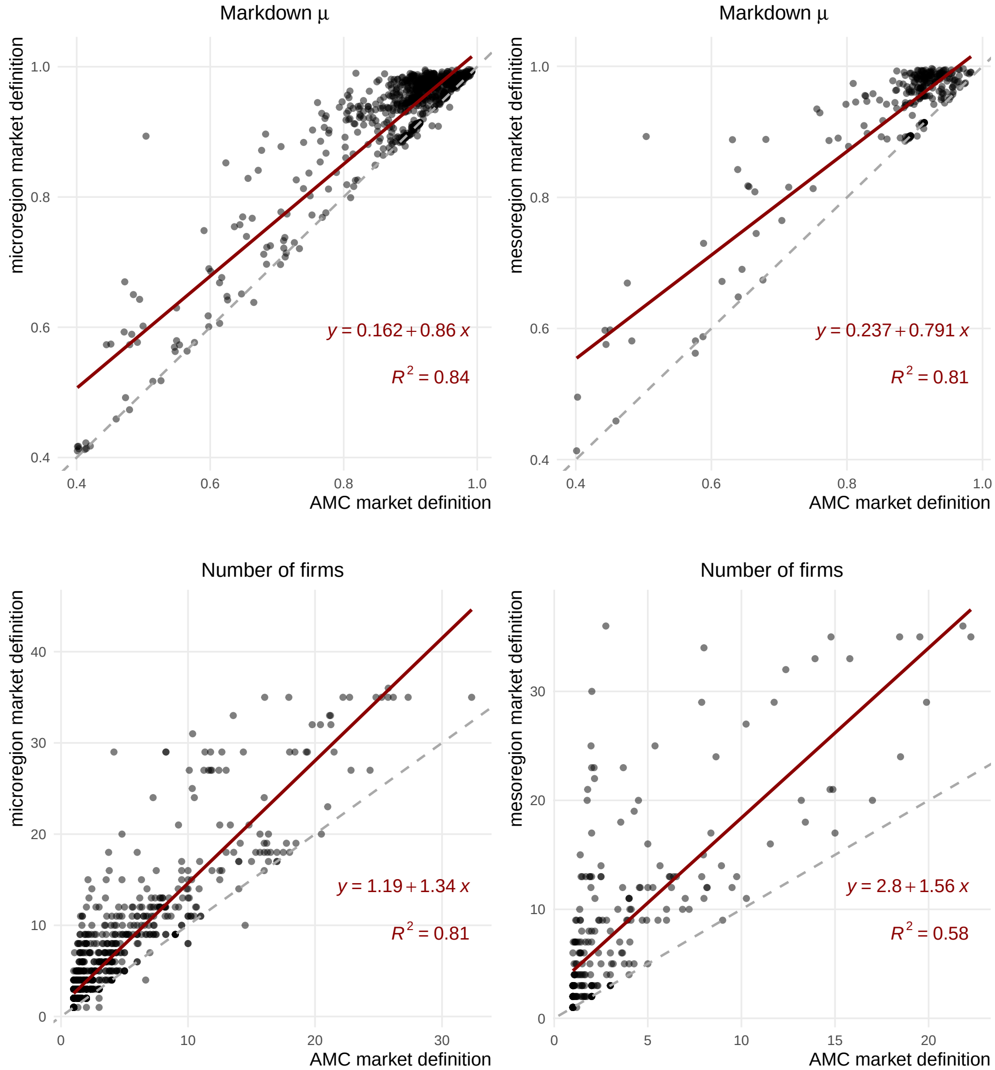

Notes: Values are pooled across all commodities for which agribusiness data is available (beef, maize, soybean). To compare AMC-level variables to larger spatial units, the AMC-level variables are averaged up to the larger unit.

Figure 3: Correlation of markdowns with emissions intensity across upstream market definitions.

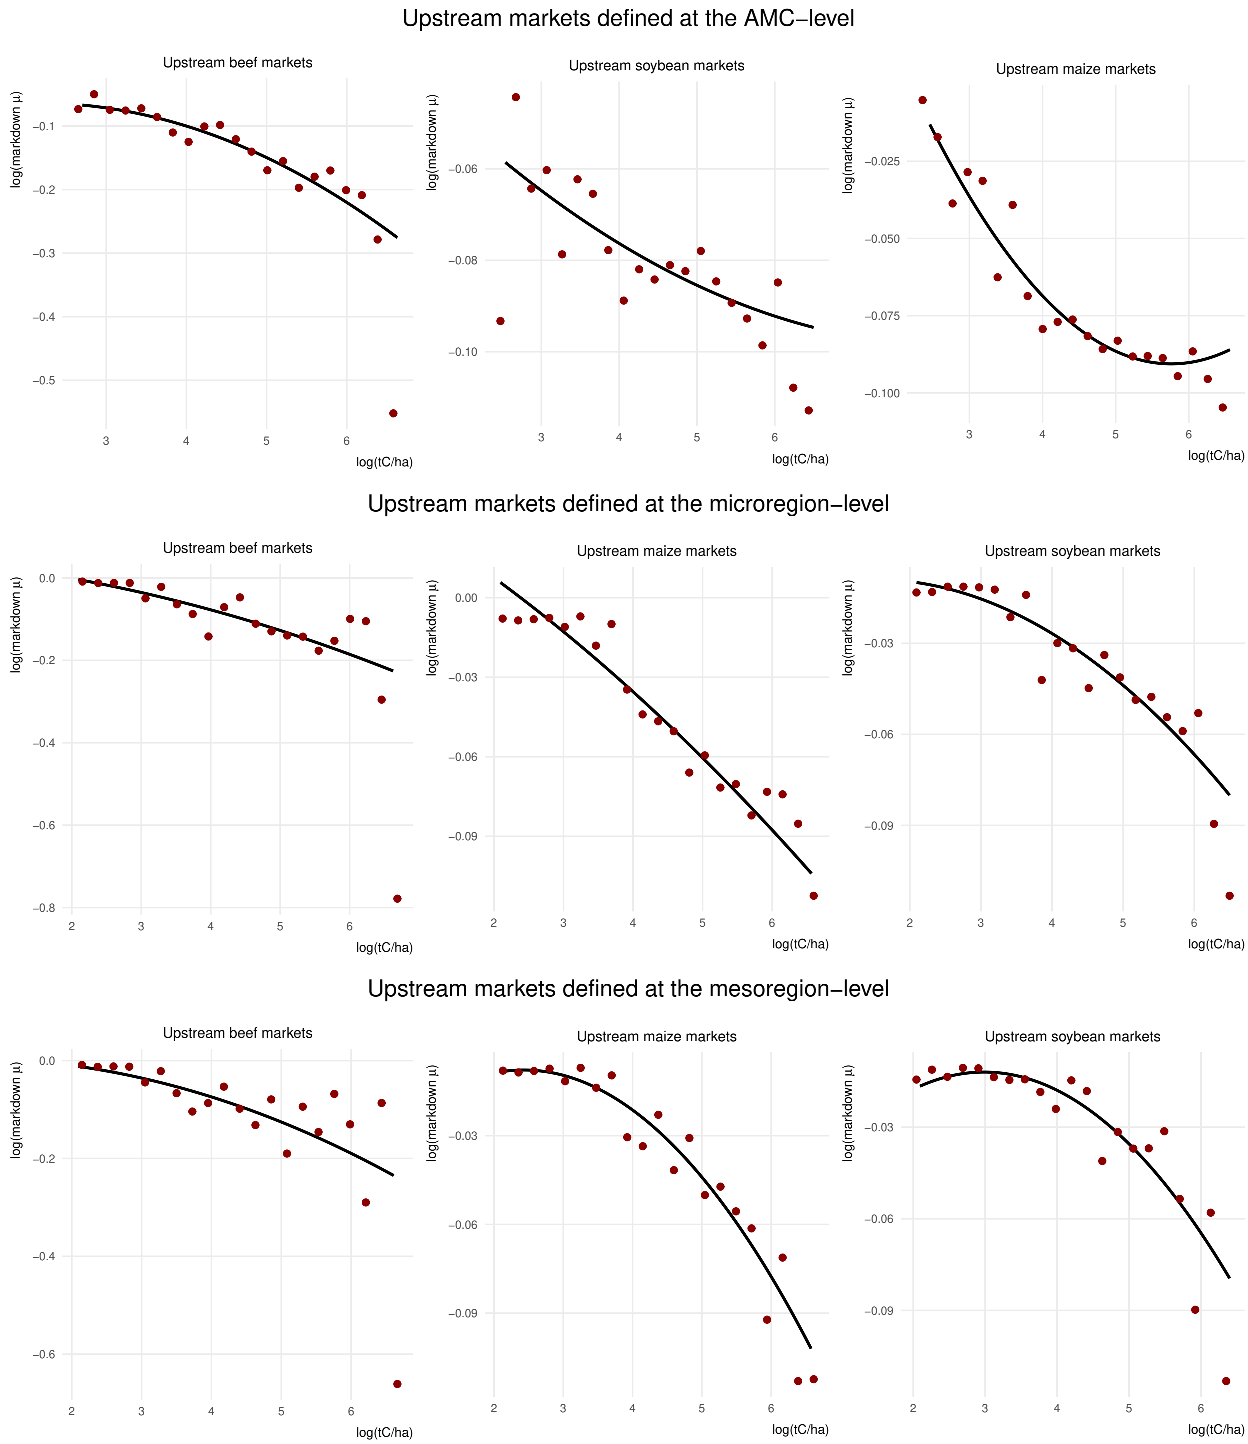

Notes: all scatter plots are constructed from upstream market-level observations on markdowns, market access, and carbon intensity. Black curves are fitted second-order polynomials of the upstream market-level observations, while red markers display the binned scatter plot of such observations. The horizontal axis is the carbon intensity of the upstream market, measured as average tonnes of carbon per hectare.

C.3.2 Correlation of markdowns with remoteness and emissions intensity

Figure 4 shows the spatial correlation between markdowns, remoteness, and emissions intensities for various commodities.

Figure 4: Correlation of model-implied markdowns with county-level observables.

A. Correlation with remoteness (measured as market access).

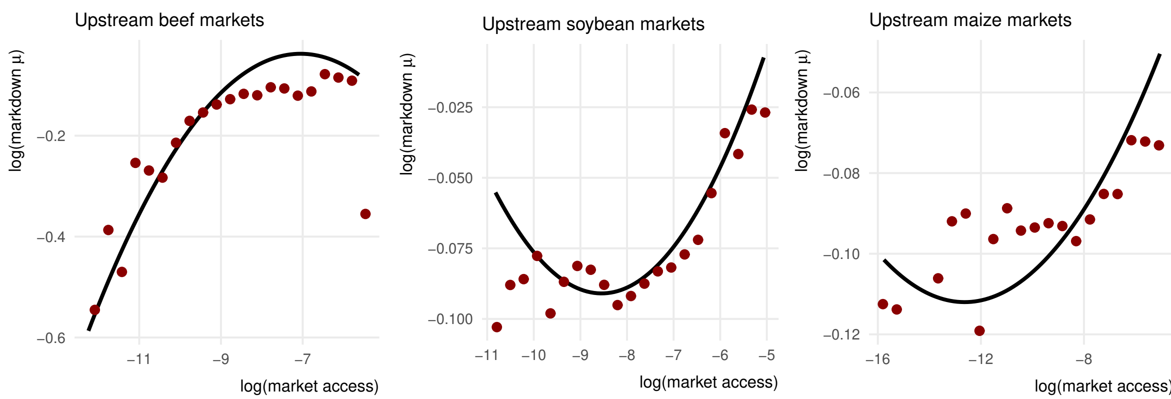

B. Correlation with upstream emissions intensity (measured as carbon stock/hectare).

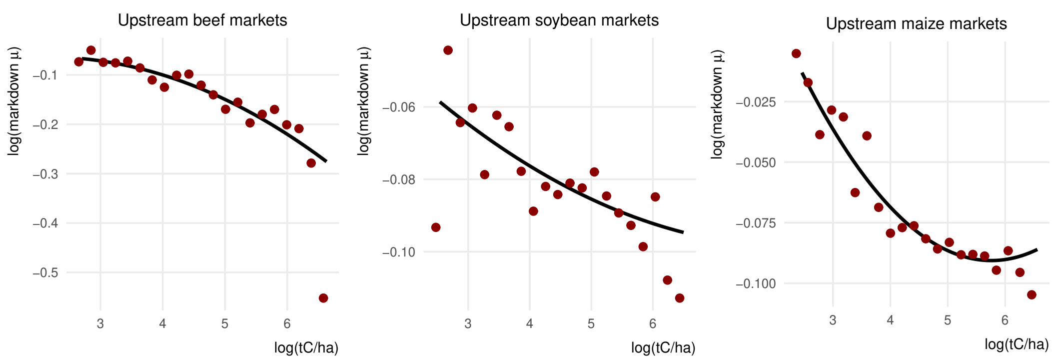

Notes: scatter plots are constructed from county-level observations on markdowns, market access, and carbon intensity. Black curves are fitted polynomials of county-level observations. Red markers are binned scatter plots. Markdowns are computed using 4. Market access is constructed as weighted-average distance from logistics hubs: $\sum_j s_j d_{ij}^{-1}$ where d_{ij} is county i-hub j distance and weights s_j reflect hub trade volume. Carbon intensity of each county is measured as average tonnes of carbon/ha.

C.3.3 Validation of model-implied markdowns to external measures of markdowns

I construct external accounting-based markdowns for the subset of commodities for which we have the data to do so (soybean, beef, maize) and compare them to model-implied markdowns. First, note the external and model-implied markdowns are not directly comparable because they are theoretically different objects. The external markdown is computed on an accounting basis as the ratio between the upstream farm-gate price the agribusiness firm pays and the downstream price it receives at the logistics hub (ports, as well as inland consignment centers). The gap between these two prices is driven by (i) agribusiness market power, but also by (ii) transport costs from the farm to the logistics hub, and (iii) any value-added the agribusiness contributes along the supply chain. By contrast, the model-implied markdown μ_i^c is a measure of pure agribusiness market power that has already netted-out transport costs and value-added. Hence, we need to control for domestic transport costs and value-added to isolate the variation in external markdowns that is driven by buyer market power, and then compare that variation to that of the model-implied markdown. To control for transport costs I use distance between upstream location and logistics hub, which is reported in the TRASE data. To control for value-added I exploit another useful feature of the data: the format in which the commodity is exported is reported. Thus, I can control for differential value-added across commodity-processed good pairs (e.g., raw soybeans-soybean oil versus raw soybeans-soybean meal) with flexible fixed effects.

Table 6 shows a positive and statistically significant relationship between external markdowns and model-implied ones, even when controlling for domestic transport costs and agribusiness value-added. Hence, the model-implied markdowns are largest in the same locations where the market power component of the external markdowns is largest. If we are willing to assume that external markdowns are to some degree indicative of market power, then the fact that the model-implied markdowns move in the same direction across space is reassuring. Having said this, accounting-based markdowns are not perfect measures of market power, so even if my model is correctly specified, we should not expect the model-implied markdowns to perfectly align with accounting-based ones. The main objective of this exercise is simply to verify that the model-implied markdowns at least go in the same direction across space as the external ones.

Table 6: External markdowns vs. model-implied markdowns.

	Dep. var.: log (external markdown)		
	(1)	(2)	(3)
log (model-implied markdown)	0.195***	0.171***	0.143***
	(0.019)	(0.020)	(0.022)
log (distance to hub)	-0.014***	-0.012***	-0.003*
	(0.002)	(0.002)	(0.002)
Commodity fixed effects		✓	
Commodity-by-product fixed effects			✓
Observations	28,810	28,810	28,810

Notes: the dependent variable is the external (accounting-based) markdown, defined as the ratio of the upstream farm gate price to the average downstream price. The independent variable is the model-implied markdown μ_i^c. "Distance to logistics hub" is an upstream location's average distance to its surrounding hubs. The set of commodities used for this exercise are: beef, maize, soybeans. SE clustered at the county-level.

C.4 Trade costs

In perfectly competitive settings, the gap between the downstream destination j price for a commodity sourced from upstream source i (p_{ijt}^c) and the upstream price (p_{it}^c) pins down bilateral trade costs τ_{ij} through the model's equilibrium conditions,

$$\log(p_{ijt}^c) - \log(p_{it}^c) = \log(\tau_{ij}). \tag{13}$$

Hence, one could estimate the relationship between unobserved trade costs τ_{ij} and a measure of observed distance d_{ij} by regressing the observed price gaps on the left hand side of 13 on distance. However, if intermediaries hold market power, price gaps are also driven by the wedge created by markdowns μ^c_{it}. Furthermore, if intermediaries are transforming the commodity c into a processed good k instead of trading it raw, the price gap also reflects the marginal product of the commodity to the processed good, $f'(q^c_{it};k)$. In these more general cases, equation 13 becomes,

$$\log(p^k_{ijt}) - \log(p^c_{it}) = \log(\tau_{ij}) - \log(\mu^c_{it}) - \log(m^{ck}_{it}), \tag{14}$$

where we have defined $m^{ck}_{it} \equiv f'(q^c_{it};k)$ as the marginal product of commodity input c for processed good k. Hence, m^{ck}_{it} measures how efficient firms are at obtaining physical units of processed good k from a physical unit of raw commodity c, and how this varies across upstream locations. Without further assumptions, trade costs τ_{ij}, markdowns μ^c_{it}, and marginal product m^{ck}_{it} cannot be separately identified from data on price gaps alone. The conduct assumption is therefore an identifying restriction: it pins down μ^c_{it} as a function of supply elasticities and the number of firms, allowing us to isolate one component of the price gap. This leaves us with two remaining components that need to be disentangled: marginal product and trade costs.

To pin down marginal product I exploit the fact that the agribusiness data reports the format in which a commodity sourced from county i is sold to destination j—namely whether it is sold in raw form (e.g., raw soybeans, live cattle) or processed form (e.g., soybean oil, boneless beef). First, by evaluating 14 only on price gaps of raw commodities (i.e., $k = c$) we can impose $m^{cc}_{it} = 1$ because there is no processing taking place, and therefore trade costs are identified. Second, for price gaps involving a processed format of the commodity we can estimate m^{ck}_{it} by exploiting price variation across different processed goods k and k' within a raw commodity-origin-destination c-i-j tuple. Intuitively, the difference in the downstream price of soybean oil versus soybean meal along the same origin-destination route ij differences out trade costs and raw commodity-c markdowns, thus isolating the marginal product component. Further implementation details on how I control for marginal product when measuring trade costs are in the next section C.4.1.

C.4.1 Controlling for value-added to measure trade costs

This section provides details on how I parameterize the trade costs and marginal product values that are required for model simulations. Starting from equation 14 we can derive the following,

$$\log(p^k_{ijt}) - \log(p^c_{it}) + \log(\mu^c_{it}) = \underbrace{\chi \log(d_{ij})}_{\log(\tau_{ij})} + \underbrace{FE_{ickt}.}_{-\log(m^{ck}_{it})} \tag{15}$$

The left hand side is observed since we have data on price ratios p^k_{ijt}/p^c_{it} and estimates for markdowns μ^c_{it} from the supply estimation exercise. On the right hand side of 15 we have i) unobserved trade costs τ_{ij}, which we parameterize as a function of observed distance d_{ij} with elasticity χ, and ii) unobserved marginal product, which we can control for non-parametrically with a fixed effect FE_{ickt} due to the richness of the agribusiness data.[2] From 15 we obtain an estimate for the distance

[2] Downstream prices are defined by upstream origin i, downstream destination j, and for each product k processed from a given commodity c. Consider the downstream price ratio between two goods k and k' produced from the same commodity c (e.g., k = soybean oil and k' = soybean meal produced from c = raw soybeans) and along the same origin-destination route ij, as in equation 16. The ratio of downstream prices identifies the ratio of marginal products because the origin-destination ij and commodity c terms cancel out. Since marginal product levels are identified for $k = c$ (because $m^{cc}_{it} = 1$), we can then use 16 to back out marginal product levels for $k \neq c$.

$$\frac{p^k_{ijt}}{p^{k'}_{ijt}} = \frac{m^{ck}_{it}}{m^{ck'}_{it}}. \tag{16}$$

elasticity $\hat{\chi}$, which we use to project trade costs as $\log(\hat{\tau}_{ij}) = \hat{\chi}\log(d_{ij})$ in model simulations. Then going back to 15 we use the projected trade costs to back out marginal product residually,

$$\log(\widehat{m}_{it}^{ck}) = \log(\hat{\tau}_{ij}) - \left[\log(p_{ijt}^{k}) - \log(p_{it}^{c}) + \log(\mu_{it}^{c})\right].$$

In the model simulations we don't model individual finished products, but we do take into account overall value-added by using the average marginal product of commodity c across all its final products $k = 1, 2, \ldots, K$, defined as $\overline{m}_{it}^{c} = \sum_k \widehat{m}_{it}^{ck}/K$. The main role this plays is to allow for differences in agribusiness productivity by location.

Appendix D. Simulations

D.1 Model parametrization details

Unobserved land productivity ζ_{it}^{c}. In section C.2 we had defined an unobservable term $\tilde{\zeta}_{it}^{c} \equiv (w_{Hit})^{-\frac{\gamma_H^c}{\gamma_L^c}} (w_{Mit})^{-\frac{\gamma_M^c}{\gamma_L^c}} \zeta_{it}^{c}$, which consisted of unobserved land productivity (ζ_{it}^{c}) and unobserved factor prices (w_{Hit}, w_{Mit}). Moreover, when estimating supply we calibrated $\tilde{\zeta}_{it}^{c}$ such that model-implied yields matched those observed in the data. We did not disentangle ζ_{it}^{c} from the non-land input prices (w_{Hit}, w_{Mit}) because it was not necessary for estimation. However, this decomposition is necessary for model simulations because the non-land input prices will be endogenously determined in equilibrium. Therefore, we incorporate the information on observed yields embedded in $\tilde{\zeta}_{it}^{c}$ by choosing ζ_{it}^{c} such that, given the simulation guess for w_{Mit} and w_{Hit}, the model implied $\tilde{\zeta}_{it}^{c}$ matches the one backed out from our estimation. The procedure is as follows,

- Given our calibrated $\tilde{\zeta}_{it}^{c}$ terms and guesses for w_{Mit}, w_{Hit}, we can back out the unobservable productivity terms as $\zeta_{it}^{c} = \tilde{\zeta}_{it}^{c} (w_{Hit})^{\frac{\gamma_H^c}{\gamma_L^c}} (w_{Mit})^{\frac{\gamma_M^c}{\gamma_L^c}}$. Hence, we are backing out the ζ_{it}^{c} terms such that their interaction with our guesses for factor prices delivers the $\tilde{\zeta}_{it}^{c}$ term from our estimation exercise.
- Notice $r_{it}^{c} = \left(p_{it}^{c}\right)^{\frac{1}{\gamma_L^c}} z_{it}^{c}\tilde{\zeta}_{it}^{c}$, i.e., it is invariant to the guess of w_{Mit} and w_{Hit}. Because we get the same r_{it}^{c} for any guess of w_{Mit} and w_{Hit}—and recall r_{it}^{c} is what determines land use shares—output is also invariant to the guesses,

$$Q_{it}^{c} = \tilde{r}_{it}^{c}\left(\pi_{it}^{c|\mathcal{C}}\right)^{-\frac{\lambda}{\theta}}\left(\pi_{it}^{\mathcal{C}}\right)^{-\frac{1}{\theta}} L_{it}^{c}, \quad \text{where } \tilde{r}_{it}^{c} \equiv r_{it}^{c}/p_{it}^{c}\gamma_L^{c}.$$

- Demand for labor and intermediates is not invariant to the guesses, because they depend on the guesses directly, conditional on output,

$$H_{it}^{c} = \gamma_H^{c}\frac{p_{it}^{c}Q_{it}^{c}}{w_{Hit}} \quad \text{and} \quad M_{it}^{c} = \gamma_M^{c}\frac{p_{it}^{c}Q_{it}^{c}}{w_{Mit}}. \tag{17}$$

- Hence, given any guess of w_{Mit} and w_{Hit} we back out the same land shares and output—but not factor demands—because we are forcing the ζ_{it}^{c} to deliver the same land share and output. Hence, when looping over w_{Mit} and w_{Hit} we are solving for factor market clearing conditional on an output level.

- We run this procedure above in the baseline laissez-faire model to recover ζ_{it}^c as model primitives. We then keep those ζ_{it}^c terms fixed in the policy counterfactuals.

Unobserved payoff to nature $A_{it}^{\mathcal{N}}$. Given the $\frac{\theta}{\lambda}$ and λ estimates, we calibrate $A_{it}^{\mathcal{N}}$ to fit the data when running our across-nest estimation.

$$\ln\left(A_{it}^{\mathcal{N}}\right) = \frac{\lambda}{\theta}\ln\left(P_{it}^{\mathcal{C}}\right) - \frac{1}{\theta}\ln\left(\frac{\pi_{it}^{\mathcal{C}}}{\pi_{it}^{\mathcal{N}}}\right). \tag{18}$$

Remaining parameters. The remaining parameters are either directly observed in the data or taken from the literature: the Cobb-Douglas factor shares, the across-sector labor supply elasticity ψ, the non-agricultural wages $\underline{w}_{Hit}$, and the imported intermediate input price w_{Mt}.

Table 7: Model parameterization details.

Model section	Parameter	Source
Supply	Land use elasticities ($\frac{\theta}{\lambda}$, λ)	Substitution elasticities from supply estimation
	Land endowment ($\bar{L}_i$)	Surface area from agricultural censuses (Brazil 2017, Argentina 2018)
	Labor endowment ($\overline{H}_{Ni}$)	Economically active population (Brazil 2010 population census, Argentina RENAPER 2023)
	Wage in non-agriculture ($\underline{w}_{Hit}$)	County-averages from Brazil 2010 population census, Argentina 2014 MECON labor survey
	World intermediate input price (w_{Mt})	Imported fertilizer prices (FAOSTAT)
	Observed land productivity (z_{it}^c)	Agronomic yields (FAO-GAEZ)
	Unobserved land productivity (ζ_{it}^c)	Residuals from within-nest estimation (described in section D.1)
	Unobserved payoff to nature ($A_{it}^{\mathcal{N}}$)	Residuals from across-nest estimation (described in section D.1)
	Cobb-Douglas parameters (γ_L^c, γ_H^c, γ_M^c)	Averaged across related literature (Sotelo, 2020; Pellegrina, 2022)
	Across-sector labor supply elasticity (ψ)	Related literature (Farrokhi and Pellegrina, 2023)
Demand	CES elasticities (η_l, η_m, η_u)	Demand estimation
	Taste shifters (a_{ijt}^c, a_{njt}^c, a_{jt}^c)	Residuals from demand estimation
	Total expenditure (X_{jt})	Expenditure on agricultural imports (FAOSTAT)
Intermediaries	Bilateral trade costs (τ_{ij})	Trade cost estimation (described in section C.4)
	Marginal product ($f'(q_i^c, k)$)	Trade cost estimation residuals (described in section C.4)
	Number of firms (N_i^c)	TRASE and DA-MAGYP
Other	LUC emissions (e_i^{LUC})	Above- and below-ground carbon stock maps (Spawn and Gibbs, 2020)
	NLUC emissions ($e_i^{c,NLUC}$)	Life-cycle GHG values (Poore and Nemecek, 2018)

D.2 Equilibrium solver algorithm

To solve the equilibrium, I use a simple iterative algorithm that exploits the monotonicity of the excess demand function in both output and factor markets. The algorithm iterates over the farm-gate price vector $\{p_i^c\}_{i,c}$ and wage vector $\{w_{Hi}\}_i$, adjusting the price and wage "guesses" based on the resulting excess demand. Note the intermediate input price vector $\{w_{Mi}\}_i$ is exogenous since supply is foreign and perfectly elastic, so we do not need to solve for market clearing in the intermediate input market. The steps of the algorithm are:

1. **Outer loop (output market).** Let the current guess for the price vector be $\{\dot{p}_i^c\}_{i,c}$.
2. **Inner loop (labor market)**. Given $\{\dot{p}_i^c\}_{i,c}$, find the wage vector $\{w_{Hi}\}_i$ that clears labor markets by using the bisection method.
   (a) The excess labor demand function at an initial guess $\{\dot{w}_{Hi}^c\}$ and $\{\dot{p}_i^c\}_{i,c}$ is,
   $$ED_i^H(\dot{w}_{Hi}) \equiv \sum_c H_i^c(\dot{w}_{Hi}; w_{Mi}, \{\dot{p}_i^c\}) - s_{H|i}(\dot{w}_{Hi}; \underline{w}_{Hi})^{(\psi-1)/\psi}\,\overline{H}_i.$$
   Do the same for another guess $\ddot{w}_{Hi}$. Pick the starting guesses $\{\dot{w}_{Hi}, \ddot{w}_{Hi}\}$ such that

$ED_i^H(\dot{w}_{Hi})$ and $ED_i^H(\ddot{w}_{Hi})$ have different signs.

(b) Compute the midpoint guess as $\widehat{w}_{Hi} = \frac{\dot{w}_{Hi}+\ddot{w}_{Hi}}{2}$.

(c) If the sign of $ED_i^H(\widehat{w}_{Hi})$ is the same as $ED_i^H(\dot{w}_{Hi})$, then update $\dot{w}_{Hi} = \widehat{w}_{Hi}$. Conversely, if the sign of $ED_i^H(\ddot{w}_{Hi})$ is the same as $ED_i^H(\ddot{w}_{Hi})$, then update $\ddot{w}_{Hi} = \widehat{w}_{Hi}$.

(d) Repeat steps (a)-(b) until $|\dot{w}_{Hi} - \ddot{w}_{Hi}| < \epsilon_H \ \forall i$, where ϵ_H is the tolerance level.

3. Having solved the labor market equilibrium on the output price guess $\{\dot{p}_i^c\}_{i,c}$, evaluate the excess demand function for each commodity c, location i,

$$ED_i^c(\dot{p}_i^c) \equiv \sum_j C_{ij}^c(\dot{p}_{ij}^c)\tau_{ij}^c - Q_i^c(\dot{p}_i^c) \quad \text{where} \quad \dot{p}_{ij}^c = \frac{\dot{p}_i^c}{\mu_i^c(\dot{p}_i^c)}\tau_{ij}^c.$$

4. Let $\{\ddot{p}_i^c\}$ be the updated guess, obtained by using $ED_i^c(\dot{p}_i^c)$ and a step size of $\alpha_Q > 0$,

$$\ddot{p}_i^c = \dot{p}_i^c + \alpha_Q ED_i^c(\dot{p}_i^c).$$

Therefore, if $ED_i^c(\dot{p}_i^c) > 0$ then $\ddot{p}_i^c > \dot{p}_i^c$, and vice-versa. Because of monotonicity of the excess demand function, the updated guess will take excess demand closer to zero. Set the updated guess as the current guess, i.e., $\{\dot{p}_i^c\} = \{\ddot{p}_i^c\}$.

5. Repeat steps 1-4 until $ED_i^c(\dot{p}_i^c) < \epsilon_Q \ \forall i, c$, where ϵ_Q is the tolerance level.

D.3 Model fit

Given the parametrization from Table 7, the model equilibrium is solved using the algorithm from Appendix D.2. Figure 5 compares model outcomes to their counterparts in the data. Note that despite using an inversion approach to discipline some parameters, this does not force a perfect fit of the model. The reason is that productivities are inverted solely from the supply estimation exercise, without imposing demand side information or solving the model.

Figure 5: Model fit (spatial distribution of acreage and output).

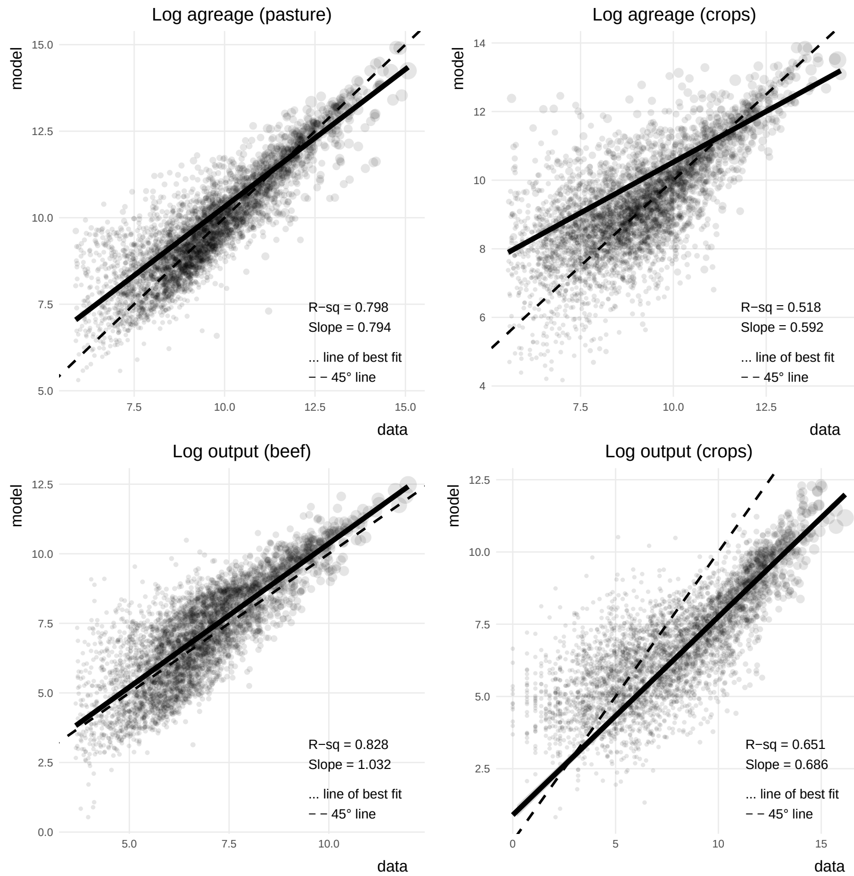

Notes: figure plots county-level model-implied outcomes versus their observed value in the data.

D.4 Welfare accounting

We begin by defining the welfare function that a (global) social planner would maximize in the model. It consists of four components: consumer surplus (from every consumer country), producer surplus (of landowners, farmers and intermediaries in South America), government revenue, and the damages from the environmental externality,

$$W = \underbrace{CS + PS + G}_{\equiv S} - E \times SCC, \tag{19}$$

where we have defined S as the surplus generated by the market, without taking into account the environmental damage.

D.4.1 Individual welfare components

To be fully general, let t_{ij}^{c} denote any downstream tax on consumers, i.e., on their bilateral consumption of commodity c, C_{ij}^{c}. Denote t_i^c as any upstream taxes on producers of c in location i, i.e., on their output Q_i^c. Finally, denote $s_i^{\mathcal{N}}$ as any upstream subsidy on landowners in location i on their forested land $L_i^{\mathcal{N}}$. On the producer side, we separate output taxes from land subsidies to separately target on-farm flow emissions (which is done by targeting a type of commodity) from land use change emissions (which is done by targeting a location, regardless of which specific commodity is produced). We now move on to specify each of the welfare components.

Environmental damages. The environmental damages are simply the sum of on-farm CO_2-eq emissions for each commodity, E_i^c, and land use change CO_2-eq emissions $E_i^{\mathcal{C}}$, evaluated at the social cost of carbon. Total emissions are,

$$E = \left(\sum_i \left(\sum_c E_i^c \right) + E_i^{\mathcal{C}} \right).$$

Recall we aren't interested in the level of W or the level of each of its components, but in how they change between a baseline equilibrium and a counterfactual one. We denote the baseline equilibrium with a superscript * and the counterfactual with a superscript **, and the difference in an outcome X as $\Delta X = X^{**} - X^*$. The welfare change that is attributed to changes in environmental damages is straightforward to compute as $SCC \times \Delta E$ where,

$$\Delta E = E^{**} - E^*.$$

Government surplus. The global government surplus is the sum of the government surpluses of each destination j,

$$G = \underbrace{\sum_i \left(\sum_c t_i^c Q_i^c + t_{i,SA}^c C_{i,SA}^c \right) - s_i^{\mathcal{N}} L_i^{\mathcal{N}}}_{G_{SA}} + \sum_{j \neq SA} \underbrace{\sum_{i,c} t_{ij}^c C_{ij}^c}_{G_j},$$

where the South American surplus G_{SA} consists of tax revenues from upstream output taxes t_i^c and downstream consumption taxes $t_{i,SA}^c$ minus subsidy expenditures on forested land $s_i^{\mathcal{N}}$. The government surplus of destinations other than South America just consist of consumption taxes, since we are assuming foreign governments cannot levy taxes or subsidies directly on upstream South American producers. The welfare change that arises from changes in government surplus (and of each destination j) is,

$$\Delta G = G^{**} - G^* \quad \text{and} \quad \Delta G_j = G_j^{**} - G_j^*.$$

Producer surplus. Producer surplus is the sum of producer surplus of intermediaries and farmers. The surplus of the intermediaries is simply their profits,

$$PS_I = \sum_{i,c} N_i^c \left[\bar{p}_i^c f_c(q_i^c) - p_i^c q_i^c \right],$$

where q_i^c is the quantity purchased by each intermediary firm and N_i^c is the number of firms. The welfare change that arises from changes in intermediary surplus is,

$$\Delta PS_I = PS_I^{**} - PS_I^*.$$

The producer surplus of farmers is captured by the farmer's after-tax and after-subsidy revenue,

$$PS_{F,R} = \sum_i \left(\sum_c (p_i^c - t_i^c) Q_i^c \right) + s_i^{\mathcal{N}} L_i^{\mathcal{N}}.$$

Notice the above captures the sum of payments to all factors (land, labor, and intermediates). The welfare change that arises from changes in farmer surplus is,

$$\Delta PS_F = PS_{F,R}^{**} - PS_{F,R}^{*}.$$

The welfare change that arises from changes in total producer surplus is,

$$\Delta PS = \Delta PS_I + \Delta PS_F.$$

Consumer surplus. Let $\mathcal{P}_{ij}^c(C_{ij}^c)$ be the inverse demand curve of consumers in j for commodity c produced in i. Then, consumer surplus is the area under the demand curve, minus the (after-tax) cost of consumption,

$$CS = \sum_j \underbrace{\sum_{i,c} \left(\int_0^{C_{ij}^c} \mathcal{P}_{ij}^c(x) dx - (p_{ij}^c + t_{ij}^c) C_{ij}^c \right)}_{CS_j},$$

where CS_j denotes the consumer surplus in destination market j (across all origin locations and commodities). We don't have a closed form for the inverse demand curve $\mathcal{P}_{ij}^c(C_{ij}^c)$, so the CS level cannot be computed directly. However, we do not need the level of CS, just the changes between counterfactuals. Consider the indirect utility of the consumer in destination j,

$$V_j \quad = \quad \max_{C_{ij}^c} \quad U_j \quad s.t. \quad \sum_c \sum_i (p_{ij}^c + t_{ij}^c) C_{ij}^c \leq X_j.$$

Note that we have a closed form for the indirect utility because of the CES setup. Given a destination j price index P_j, we have $V_j = X_j / P_j \implies X_j = V_j P_j$. Hence, to attain utility level V_j at P_j we would need X_j of income. Consider a baseline and counterfactual equilibrium denoted by * and **, which yield indirect utilities of V_j^* and V_j^{**} at price indices P_j^* and P_j^{**}. Define x_j as how much more we have to increase income from the baseline so that the counterfactual level of utility V_j^{**} can be attained at the baseline prices P_j^*,

$$V_j^{**} = (X_j^* + x_j) / P_j^* \implies x_j = P_j^* V_j^{**} - X_j^*.$$

We can therefore evaluate x_j from the model counterfactuals, and then use it to compute the change in total consumer surplus as follows,

$$\Delta CS = \sum_j \Delta CS_j \quad \text{where} \quad \Delta CS_j = x_j.$$

D.5 Additional simulation results

D.5.1 Regressivity results

Figure 6 shows that the rate at which farmer income is effectively taxed (as a result of the downstream tax) is lower in higher-income locations. The left panel shows the impact on farm-gate prices expressed as percentage point declines, i.e., the policy's implied ad-valorem tax on farm-gate prices. The right panel shows the impact on farmer income expressed as percentage point declines, i.e., the policy's implied income tax on farmers. Both measures indicate the implied income tax is highest for farmers in the poorest regions, making the policy regressive across space.

Figure 6: Regressivity of downstream taxes.

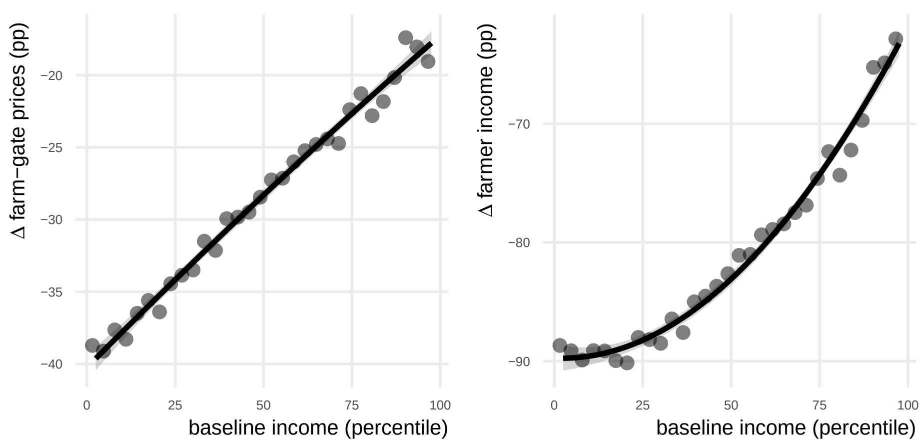

Notes: all scatter plots are constructed from county-level observations. Black curves are fitted second-order polynomials of county-level observations, while dark markers display the binned scatter plot. The vertical axes is the percentage point change in (i) the upstream farm-gate price and (ii) farmer income, as a result of the downstream tax. The horizontal axes show the baseline income of the county, expressed in percentile terms. A positive slope of the scatter plots means that counties that are higher-income at baseline have their income effectively taxed at lower rates.

D.5.2 Leakage results

Figure 7 compares the abatement impact of the downstream tax when imposed by the domestic regulator versus a foreign trade partner. In the latter case, regulation is incomplete across consumer markets, hence the drop in exports to the EU is offset in equilibrium by increased consumption in non-EU markets, including the domestic market.

Panel A shows the tax's corrective potential is substantially limited by this consumption leakage effect: over 80% of the emissions reductions attributed to the drop in EU consumption are offset by increased consumption elsewhere. Panel B shows the case with complete regulation as contrast: there is no re-routing of shipments, so emissions drop across all destination markets. Note the emissions changes in the figure take into account substitution away from beef and into crops, i.e., domestic production leakage from beef to crops (given beef is taxed at a higher rate than crops due to its higher emissions footprint). Specifically, the emission changes reported in the figure include the growth in crop-related emissions generated by substitution from beef into crops. Therefore, the fact that all bars in panel B are negative indicate the drop in emissions from beef production more than offsets any increase in emissions from crop production.

Figure 7: Emissions abatement potential of downstream taxes.

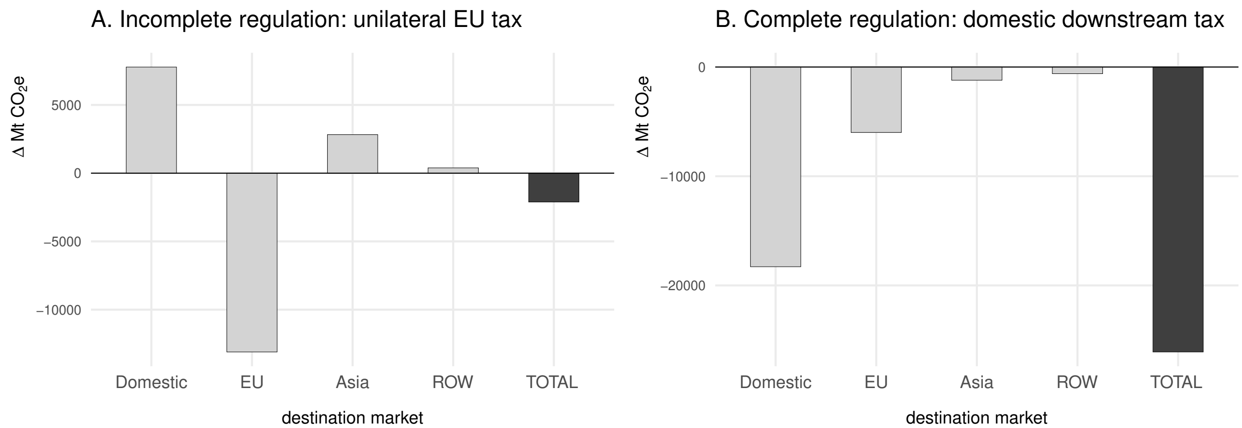

Notes: emissions are computed across all commodities, which are then aggregated up to the destination level, thus delivering the changes in emissions attributed to the changes in consumption of each destination (shown in grey, with the net emissions impact across all destinations in black).

D.5.3 Extension with migration

Related work in the agricultural trade literature typically models labor as perfectly mobile within a country, but not across countries. If workers do not have idiosyncratic shocks over locations within a country, this implies a single country-level wage in equilibrium, as in baseline specifications of Costinot, Donaldson and Smith (2016) and Farrokhi and Pellegrina (2023). To allow wages to vary by both sector (agriculture, non-agriculture) and location, I follow an extension from Farrokhi and Pellegrina (2023) that uses a Roy model of labor supply where workers have idiosyncratic shocks over sectors and locations. Let $s_{H|it}$ denote the share of workers choosing agriculture conditional on choosing location i, and let s_{it} denote the share of nation-wide population choosing location i. Assuming the idiosyncratic shocks follow an extreme value distribution, these shares have the following closed form expression,

$$s_{H|it} = \frac{w_{Hit}^{\psi}}{w_{Hit}^{\psi} + \underline{w}_{Hit}^{\psi}}, \quad s_{it} = \frac{V_{it}^{\psi}}{\sum_{i'} \left(V_{i't}^{\psi}\right)}, \quad V_{it} \equiv \nu_{it} \left(w_{Hit}^{\psi} + \underline{w}_{Hit}^{\psi}\right)^{\frac{1}{\psi}},$$

where ψ is the elasticity of substitution across sectors and locations and ν_{it} is a location shifter. Note that in the baseline model with no migration the s_{it} term is exogenous and taken from the data as the share of nation-wide population in county i. The share of nation-wide population working in location i and in agriculture is $s_{Hit} = s_{H|it} \times s_{it}$. Effective labor supply to location i is $s_{Hit}^{(\psi-1)/\psi}\overline{H}_t$, where $\overline{H}_t$ is nation-wide population. Note $(\psi-1)/\psi < 1$ reflects selection due to heterogeneity (without idiosyncratic shocks the effective labor supply would simply be $s_{Hit} \times \overline{H}_t$).

To make my migration extension as closely comparable to my baseline specification, I calibrate the ν_{it} parameters so that the model-implied population shares s_{it} match the population shares in the data (which are the same as the population shares in the baseline model). Given location i population shares from the data, denoted s_{it}^{*}, and the equilibrium wages $\{w_{Hit}^{*}, \underline{w}_{Hit}^{*}\}$ from the baseline model we can find the vector ν_{it} such that the following holds,

$$s^*_{it} = \frac{\nu_{it}^{\psi}\tilde{V}_{it}^{\psi}}{\sum_{i'}\left(\nu_{i't}^{\psi}\tilde{V}_{i't}^{\psi}\right)} \quad \text{where } \tilde{V}_{it} = \left[(w^*_{Hit})^{\psi} + (\underline{w}^*_{Hit})^{\psi}\right]^{\frac{1}{\psi}}.$$

This guarantees that the population distribution of the baseline model and the migration extension coincide in the laissez-faire equilibrium. The only modification to the equilibrium solver algorithm from section D.2 is that the labor supply function changes because the population shares are endogenously determined. At a given wage and price guess $\{\dot{w}^c_{Hi}\}$ and $\{\dot{p}^c_i\}_{i,c}$, the excess labor demand function is now as follows,

$$ED^H_i(\dot{w}_{Hi}) \equiv \sum_c H^c_i(\dot{w}_{Hi}; w_{Mi}, \{\dot{p}^c_i\}) - s_{Hi}\left(\dot{w}_{Hi}; \underline{w}_{Hi}\right)^{(\psi-1)/\psi}\overline{H}.$$

Implications for policy counterfactuals. To assess the role of migration for the paper's core results on policy targeting, I compare my downstream tax counterfactual results with and without migration. Figure 8 is a version of Figure 4-A from the main text, which shows the change in upstream production in the baseline specification without migration (in solid lines) and the migration extension (in dashed lines). The downstream tax leads to a larger drop in core emissions-light areas relative to frontier emission-intense areas, both when there is migration and when there isn't. This is the paper's main spatial mistargeting result, and it is reflected by the positives slopes of the figure under both mobility assumptions.

Figure 8: Role of migration for tempering mistargeting of a downstream tax.

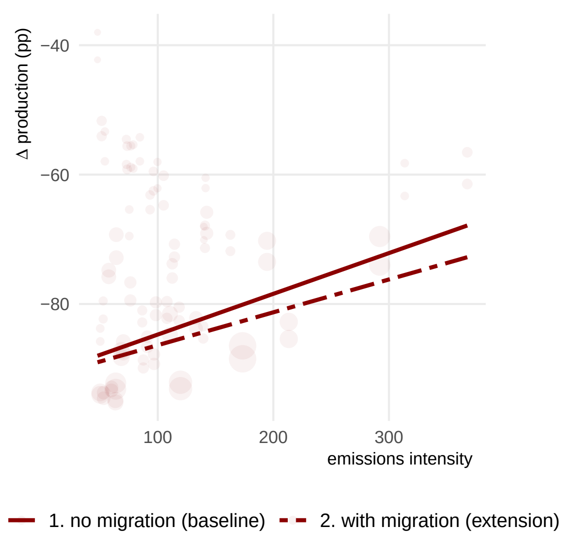

Notes: the figure plots the relationship between upstream emissions intensities and production responses to a downstream tax under the default conduct assumption of imperfect competition (for the beef sector). The horizontal axis reports the emissions intensity of each upstream location (tonnes of CO_2e/tonne of output). The vertical axis reports the percentage point change in upstream production due to the downstream tax. Lines of best fit are reported separately for the baseline assumption of no migration (solid line) and the extension with migration (dashed line).

When migration is introduced, the slope of the figure remains positive but is flatter. This means that mistargeting remains, but is quantitatively less severe. Hence, allowing migration tempers some of the spatial mistargeting, but doesn't fully eliminate it. To see why, notice the downstream tax is a larger tax on beef than on crops because of beef's larger carbon footprint. Because beef production is dominant in frontier areas (where it has comparative advantage), the downstream tax ends up taxing the good in which the frontier region specializes, thus inducing labor to migrate away to core regions. For this reason, production ends up dropping relatively more in frontier areas than in core areas when migration is allowed, thus explaining the flattening of the slope.

Appendix E. Institutional background

E.1 Supply chain structure

E.1.1 Beef cattle

Commercial cattle rearing consists of three stages: breeding, backgrounding, and finishing. In the breeding stage, calves are kept with their mothers, feeding exclusively on milk and grass. Backgrounding follows, which consists of developing the calf's frame at a moderate pace so that it resists the rapid weight gains it will be subject to in the final finishing stage. Breeding takes place on pasture, whereas diets in the backgrounding and finishing stage typically include feed. If the finishing is done at a feedlot then diets are exclusively feed-based. Although instances of vertical integration between slaughterhouses and feedlots certainly exist, integration further upstream with ranchers at the breeding stage is rare, resulting in a fragmented supply chain. Hence, the upstream actors most likely to be involved in deforestation practices—the breeders, since they rely exclusively on pasture to feed their cattle—are typically not integrated with the downstream agribusiness firms. This poses certification challenges that are unique to cattle, in contrast to commodities with shorter and more vertically integrated supply chains, such as coffee.

Vertical integration between ranchers and meatpackers is less common in beef than in other livestock sectors such as poultry for structural reasons. Commercial poultry farming involves high start-up infrastructure costs to build chicken houses, and requires a steady supply of grain-based feed and antibiotics to prevent disease outbreaks. In a typical vertical arrangement, downstream processors provide the initial capital to build the infrastructure, and may also finance the inputs, in exchange for a share of the farmer's output. By contrast, cattle require grazing on extensive pasture land in the early stages of their life, therefore infrastructure and feed costs are not critical constraints, and they take at least two years to reach slaughter weight. Both of these factors result in limited scale economies from integration (Crespi and Saitone, 2018). Moving beyond beef cattle, vertical integration is substantial in sectors such as fresh fruits and vegetables because of the wide variation in quality and their high perishability: integration allows processors to control quality and minimize lags between harvest and processing (Otsuka, Nakano and Takahashi, 2016). By contrast, the gains from vertical integration in cash crops such as soybeans are smaller due to their storability and relatively homogenous quality.

In Argentina, the beef supply chain begins with a primary sector consisting of approximately 130,000 breeding and finishing establishments, of which 52% hold less than 100 head of cattle.[3] Live cattle are purchased from these establishments by slaughterhouses, of which there are approximately 400. Vertical integration between breeders and slaughterhouses is limited compared to other animal agriculture sectors such as poultry. Around half of cattle transactions take place directly between the breeding establishment and the slaughterhouse, a quarter is intermediated by

[3]Source: SENASA.

a broker, and the remaining quarter takes place in local spot markets/auctions (Bisang, Santagelo, Anlló and Campi, 2007). In Brazil, as of 2017 there are 2.5 million cattle ranching establishments, 73% of which are less than 50 ha, 79% of which hold less than 50 head of cattle, and 76% of which are family farms.[4] As in Argentina, the supply chain is vertically unintegrated and divided into three parts: breeders, finishers, and slaughterhouses. Note that over 70% of beef output in both countries is consumed domestically.[5] This is in contrast to cash crops such as soybeans, which are almost entirely exported. Argentina and Brazil together hold approximately 20% of the world's cattle stock, produce 20% of the world's beef, and account for 30% of world beef exports.[6]

E.1.2 Soybean

In Argentina, there are roughly 60,000 soybean producers of heterogenous size. 80% of producers are relatively small, with a planted area less than 260 ha.[7] Raw soybeans are purchased from farmers by agribusiness firms such as Cargill, Louis Dreyfus, and ADM, either for direct export in raw format, or for processing into soybean oil and meal. A quarter of total soybean production is exported raw, while the remainder is processed, with 70-80% of the processed products exported. As of 2012, there are a total of 50 soybean crushing plants, corresponding to 37 firms. The top 8 firms account for 80% of national processing capacity and over 90% of exports (Pierri, Wesz Junior et al., 2017). These firms typically hold state of the art storage and processing facilities (grain elevators, silos, etc.) as well as their own port terminals along the Paraná river. High entry costs exist in the form of large investments in physical infrastructure and the development of a logistical/marketing network. The industry structure is similar in Brazil, the main difference being that export of raw soybeans is more common (Burgos, Mattos and Medina, 2014). As of 2017 there are 236,245 soybean farms—of which 84% are less than 200 ha and 70% are family farms—and 96 crushing facilities in Brazil.[8]

Both countries underwent similar liberalization and privatization processes in the 1990s that spurred agricultural commodity exports and the expansion of multinational agribusiness firms. First, price support programs were either eliminated or scaled down in both countries. For example, Argentina dismantled its *Junta Nacional de Granos* in 1991, whose purpose was to purchase grain crops at a guaranteed minimum price. Second, exports duties on agricultural commodities were either eliminated or significantly reduced, raising the international competitiveness of local producers. For example, the 1996 Kandir Law in Brazil exempted export goods from paying domestic circulation duties. Third, the management of port authorities was decentralized to subnational/local governments, establishing the legal framework for large multinationals to build their own processing facilities within private port terminals. In Argentina this was enabled with the approval of Law 24,093 in 1992, and in Brazil with Law 8,630 in 1993.

The main driver of soybean demand throughout the sample period has been Chinese economic growth, reflected by the increase in Chinese soybean imports, in particular from Argentina and

[4]Sources: IBGE Table 6783 and Table 6910 and Table 6783.

[5]Of the beef that is exported, over half goes to Asia, and 20-30% to Europe. Jointly, Argentina and Brazilian beef exports account for nearly 30% of world beef exports. By itself, Brazil's share is 20%, making it the world's single largest beef-exporting country. Source: FAOSTAT.

[6]The other major countries with large cattle stocks are India (12% of the world cattle stock), China (7%), and the United States (6%). The next major producers after Argentina and Brazil are the United States (18% of world production) and China (9%). The next major exporters after Argentina and Brazil are the United States, Australia, and India, with world export shares between 10-15% each. Source: FAOSTAT.

[7]Source: MAGYP.

[8]Source: IBGE Table 6957, Table 6959, and TRASE.

Brazil in Figure 9. The main use of imported soybeans in China is for animal feed, particularly the hog industry. There is substantial work, both policy institution reports and academic papers, documenting the link between Chinese economic growth, shifts towards more meat-intense diets, and the ensuing demand for soybean-based animal feed. For example, see He, Baiocchi, Hubacek, Feng and Yu (2018) for Chinese shifts towards meat-intensive diets, Peine (2013) for the link between the Chinese pork industry and Brazilian soybeans, and on the policy side, soybean outlook reports from the USDA.

Figure 9: Evolution of soybean imports (China versus World).

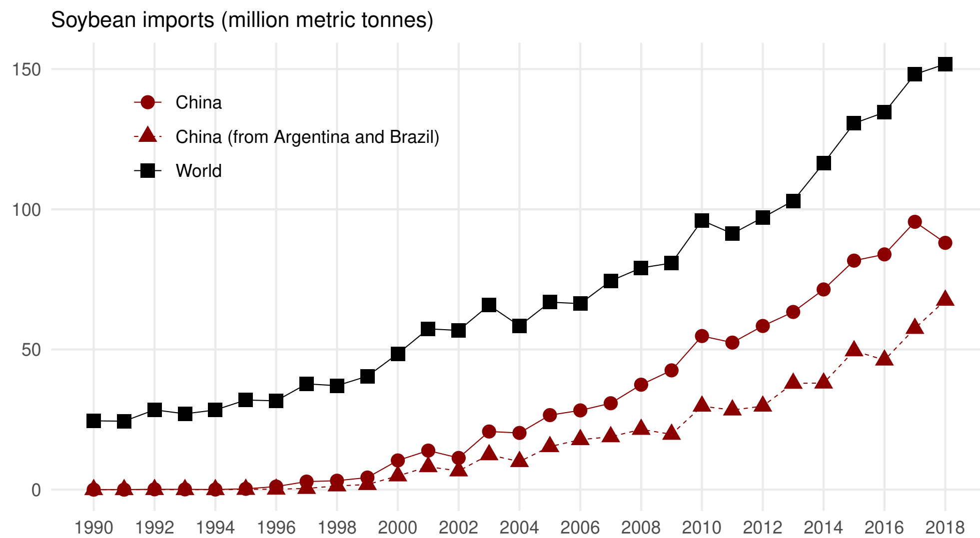

Notes: the time-series plot above is constructed using soybean import quantities (million metric tonnes, in raw bean format) from FAOSTAT. "China" corresponds to mainland China in the FAOSTAT classification.

E.1.3 Maize

The maize supply chain is similar to that of soybeans in that Argentina and Brazil consistently rank among the world's top 10 producers and exporters, with most exports going to Asia. In terms of share of global maize production, Brazil holds 7% and Argentina holds 4%. In terms of share of the global maize export market, Brazil holds 30-40% while Argentina holds 10-15%. Production is slightly less export-oriented than soybeans, with export shares of production ranging between 30-40% for Brazil and 50-60% for Argentina. A notable similarity with the soybean sector is that the same agribusiness firms tend to overlap in both commodities. For example, Bunge and Cargill are large players in both maize and soybeans.

E.1.4 Other commodities

Beyond cattle, soybean, and maize, there are other commodities we used in the estimation of the land use model: wheat, rice, sunflower, sugarcane. Given we do not have agribusiness data for these commodities, we only briefly discuss the features of their supply chains in this subsection. Note that the commodities for which we do have agribusiness data (cattle, maize, soybean) and which are the focus of the analysis, represent over 85% of agricultural land in both countries. Hence, to the extent the remaining commodities for which we do not have agribusiness data represent a small land share (wheat, rice, sunflower, sugarcane), we do not expect omitting the specifics of their market structure to play a major quantitative role. All market shares reported over the next paragraphs for each commodity are from FAOSTAT, and are averages between 1995-2018.

Rice. Brazil often ranks within the top 15 rice producers and exporters in the world, while Argentina is not a major rice producer nor exporter. In terms of shares of global production, Brazil holds 2% and Argentina holds less than 1%. In terms of shares of the global export market, Brazil holds roughly 10% while Argentina holds less than 5%. Farmers sell their output to cooperatives/traders, similarly as in maize and soybeans. In contrast to maize and soybeans, most production is consumed domestically (less than 20% of Brazilian production and less than 35% of Argentine production is exported). Traders purchasing the raw commodity from farmers primarily add value by providing storage, drying, and processing services (da Cunha and Wander, 2023).

Wheat. Argentina typically ranks among the world's top 10 producers and exporters of wheat, with roughly 60% of its output being exported. Brazil is not a major player in wheat production, with output levels being roughly half that of Argentina. In terms of shares of global wheat production, Argentina holds 2% while Brazil holds less than 1%. In terms of shares of the global wheat export market, Argentina holds 7% while Brazil holds less than 2%. Similarly to maize and soybean, intermediaries purchase wheat from farmers and add value through storage/crushing/milling services, after which the semi-processed product is sold on to food processors (Bisang, 2010).

Sunflower. Argentina typically ranks among the world's top 10 producers and exporters, while Brazil is not a major player. In terms of shares of global sunflower seed production, Argentina holds 10% while Brazil holds less than 1%. In terms of shares of the global sunflower export market, Argentina holds 5% while Brazil holds less than 1%. Upon purchasing the sunflower seed from farmers, agribusiness firms typically provide crushing services to separate sunflower oil from meal (Ingaramo and Feoli, 2008).

Sugarcane. Brazil is one of the world's largest sugar cane producers, while Argentina is a minor player. In terms of shares of global sugarcane, Brazil holds 30-40% and Argentina holds less than 2%. In terms of shares of the global cane sugar export market, Brazil holds 50% while Argentina holds less than 1%. In contrast to other crops, sugarcane is not only used for food but also as feedstock for ethanol-blend gasoline, with roughly half of the harvest used for fuel (USDA.)

E.2 Land clearing process and release of above- and below-ground carbon

Both above- and below-ground carbon is released during the deforestation process. Above-ground carbon is primarily stored in tree trunks, leaves, and other above-ground forms of biomass, while below-ground carbon is mainly stored in roots, soil, as well as in the form of peat. Below-ground carbon is released after the soil is disturbed through oxidation, and as roots/peat are left to decay or directly burned. Different forests vary in the extent to which carbon is primarily stored above- or below-ground. For example, for the Amazon, 50-60% is stored above-ground.[9] By contrast, most carbon in Indonesian rainforests, is stored below-ground in the form of peat.

The speed at which both types of carbon are released depends on how the deforestation process takes place. For example, if the vegetation is burned, the carbon is immediately released into the atmosphere. However, if vegetation is cut and left to decay, the carbon is gradually released

[9] Source: Scientific American.

as biomass decomposes. In the context of South America, given the primary driver of deforestation is to use the land for agriculture rather than to exploit the timber resource, it is common for deforestation to take place through the burning of vegetation. Apart from being a cost-efficient way to deforest, the resulting ash provides nutrients for the soil that aids conversion of land to agricultural use. Given both above-ground and below-ground biomass is disturbed when fully converting forested land to agricultural use, in my model quantification I assume both types of carbon are released into the atmosphere. Because the model counterfactuals are across steady states, there is not really a notion of speed of release of carbon in the model: I simply assume that between two counterfactuals the full stock of carbon in a deforested area is released. I have performed robustness using only above-ground carbon and the results do not significantly change because i) most carbon in the Amazon context is aboveground carbon, and more importantly ii) aboveground and belowground carbon are highly spatially correlated (as the carbon maps in the main draft show). Hence, the spatial correlation between carbon intensity and market power is the same regardless of whether we use both carbon measures or the aboveground measure alone.

E.3 Agricultural greenhouse gas emissions

Figure 10: Alternative measures of emissions footprints (per kcal and per protein content).

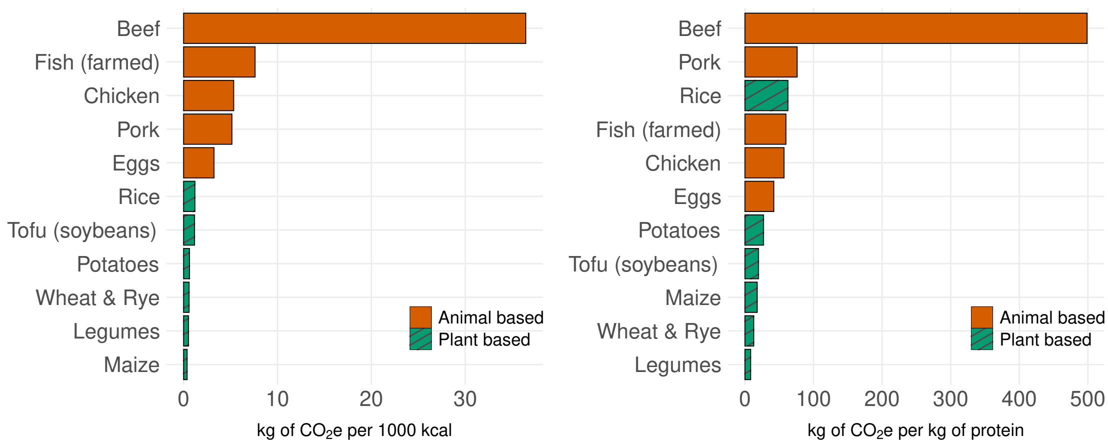

Notes: the figures above are constructed from the publicly available dataset on food emissions footprints (kg of CO_2e per 1000 kcal and per kg of protein) from Poore and Nemecek (2018).

References

Bisang, R. (2010). La Cadena de Valor del Trigo: Oportunidades y Desafíos en el largo plazo. In: *El trigo en Argentina y el mundo (19 de mayo de 2010, Buenos Aires, Argentina)*.

Bisang, R., Santagelo, F., Anlló, G. and Campi, M. (2007). Mecanismos de formación de precios en los principales subcircuitos de la cadena de ganados y carnes vacunas en la Argentina. Tech. rep., Convenio CEPAL-IPCVA.

Burgos, M., Mattos, E. and Medina, A. (2014). *La soja en Argentina (1990-2013): Cambios en la cadena de valor y nueva articulación de los actores sociales*. Centro de Economía y Finanzas para el Desarrollo de la Argentina, CEFID AR.

Costinot, A., Donaldson, D. and Smith, C. (2016). Evolving comparative advantage and the impact of climate change: Evidence from 1.7 million fields around the world. *The Journal of Political Economy* 124(1):205–248.

Crespi, J.M. and Saitone, T.L. (2018). Are Cattle Markets the Last Frontier? Vertical Coordination in Animal-Based Procurement Markets. *Annual Review of Resource Economics* 10:207–227.

da Cunha, C.A. and Wander, A.E. (2023). Price regime changes in the rice value chain in the State of São Paulo, Brazil. In: *In: CONGRESSO DA SOCIEDADE BRASILEIRA DE ECONOMIA, ADMINISTRAÇÃO E SOCIOLOGIA RURAL*.

Farrokhi, F. and Pellegrina, H.S. (2023). Trade, technology, and agricultural productivity. *Journal of Political Economy* 131(9):2509–2555.

Godar, J., Persson, U.M., Tizado, E.J. and Meyfroidt, P. (2015). Towards more accurate and policy relevant footprint analyses: tracing fine-scale socio-environmental impacts of production to consumption. *Ecological Economics* 112:25–35.

He, P., Baiocchi, G., Hubacek, K., Feng, K. and Yu, Y. (2018). The environmental impacts of rapidly changing diets and their nutritional quality in china. *Nature Sustainability* 1(3):122–127.

Ingaramo, J. and Feoli, C. (2008). La importancia del complejo girasol en la Argentina. *El cultivo de girasol en la región semiárida pampeana* 1.

Otsuka, K., Nakano, Y. and Takahashi, K. (2016). Contract farming in developed and developing countries. *Annual Review of Resource Economics* 8:353–376.

Peine, E.K. (2013). Trading on pork and beans: Agribusiness and the construction of the Brazil-China-soy-pork commodity complex. In: *The ethics and economics of agrifood competition*, pp. 193–210. Springer.

Pellegrina, H.S. (2022). Trade, productivity, and the spatial organization of agriculture: Evidence from Brazil. *Journal of Development Economics* 156:102816.

Pierri, J., Wesz Junior, V. et al. (2017). "Soyafication" in Argentina and Brazil (1980-2014): Influence of public policies, transnational companies, and the dependent economic structure. *Extensão Rural* 24(1):20–41.

Poore, J. and Nemecek, T. (2018). Reducing food's environmental impacts through producers and consumers. *Science* 360(6392):987–992.

Rajão, R., Soares-Filho, B., Nunes, F., Börner, J., Machado, L., Assis, D., Oliveira, A., Pinto, L., Ribeiro, V., Rausch, L. et al. (2020). The rotten apples of Brazil's agribusiness. *Science* 369(6501):246–248.

Sotelo, S. (2020). Domestic trade frictions and agriculture. *The Journal of Political Economy* 128(7):2690–2738.

Spawn, S. and Gibbs, H. (2020). Global Aboveground and Belowground Biomass Carbon Density Maps for the Year 2010. *ORNL DAAC* .

Train, K.E. (2009). *Discrete choice methods with simulation*. Cambridge University Press.