

Contract Design for Afforestation Programs

Nicole Immorlica^{1,2}, Wanyi Li^{3*} and Brendan Lucier¹

¹Microsoft Research

²Harvard University

³Stanford University

wanyili@stanford.com, {nicimm,brlucier}@microsoft.com

Abstract

Trees on farms provide environmental benefits to society and improve agricultural productivity for farmers. We study incentive schemes for afforestation on farms through the lens of contract theory, designing conditional cash transfer schemes that encourage farmers to sustain tree growth. We capture the tree growth process as a Markov chain whose evolution is affected by the agent’s (farmer) actions – e.g., investing costly effort or cutting the tree for firewood. The principal has imperfect information about the agent’s costs and actions taken, and wants to maximize long-run tree survival with minimal payment. We show how to calculate the optimal contract structure in our model: notably, it can involve time-varying payments and may incentivize the agent to join the program but abandon it prematurely.

1 Introduction

The UN’s Sustainable Development Goal #15 challenges society to sustainably manage forests, combat desertification, halt and reverse land degradation, and halt biodiversity loss. Afforestation and reforestation, which attempt to reverse land degradation, have recently gained attention as a clear next step towards achieving this goal [IPCC, 2019]. One opportunity for afforestation is to grow trees on agricultural land owned by farmers. Mature indigenous trees on farms not only improve biodiversity and carbon storage capacity, but also deliver robust water and soil quality which improves the long-term health of the farm. However, smallholder farmers in developing countries often either do not grow trees or abandon them before reaching maturity, because of the slow and risky nature of the growing process. This is exacerbated by farmers’ cash constraint and limited labor capacity. This raises a large set of questions regarding how best to help farmers overcome barriers and how to operationalize afforestation programs.

A typical afforestation program offers a many-year-long contract which is a payment schedule conditional on tree survival. Such programs have benefited from recent advance-

ments in machine learning and AI to process remote sensing data (satellite imagery), which have made it possible to monitor land use change cheaply [Hansen *et al.*, 2013; Jean *et al.*, 2019; Dao *et al.*, 2019; Lütjens *et al.*, 2019]. Participants get paid a percentage of the total payment at each monitoring round but the exact payment might be adjusted downwards if not all the trees have survived. For example, in [Jack, 2013], the author conducted a randomized controlled trial in Malawi where farmers were asked to grow trees over a period of 3 years. They were paid in equal installments after 6 months, then, 1, 2 and 3 years adjusted by the number of survived trees. The author observed that farmers had private information regarding their likelihood of following through to the end of the 3 year period, and that some farmers dropped out of the contract over time even though they all initially agreed to take up the task. Motivated by this experiment, we ask: if afforestation programs are designed for the long term and at a large scale, can the program designer optimize over the contract space to minimize abandonment of this slow and costly task?

Our research is a first attempt to create an analytical framework for afforestation incentive schemes using contract theory. We set up a principal-agent model, where the principal (referred to as she) can be an NGO, a local government, or any buyer of ecosystem services, and the agent (referred to as he) represents a smallholder farmer. The principal contracts with the agent to procure tree-growing services on his farm. The goal of the principal is to maximize the value of the trees grown over time minus the payments issued to the agent to incentivize this growth. There is a large literature on canonical contract design [Holmstrom and Milgrom, 1987; Ross, 1973], dynamic contract design [DeMarzo and Sannikov, 2006] and, more recently, robust contract design [Dütting *et al.*, 2019; Carroll, 2015]. A closely related work with ours is [Levin, 2003] which is also a dynamic contracting problem where both adverse selection and moral hazard may be present; further, there is two-sided limited liability. The author shows that it is sufficient to consider stationary contracts when searching for optimal dynamic contracts.

Our model’s defining feature is unique and specific to afforestation contracts, where environmental benefits depend on the state of the world and not just a stochastic outcome of the agent’s action. The dynamics of tree growth follow a Markov Chain where the state space is the tree age (or,

*Contact Author

similarly, tree height or canopy size). Natural risks and the agent's efforts both influence the steady state distribution of this Markov chain. Further, the effort choices of the agent are influenced by the natural risks, the payments from the principal, and also their private type (e.g., green-thumb or not) which determines the cost of their effort. The principal can observe the state of the tree, but does not know the agent's type nor effort. The principal cares about who *takes up* the contract, and whether and how the agent *follow through* in the long run.

We show that the optimal payment structure is simple to compute and has clear economic interpretations. For example, if tree maintenance costs are constant over time then the optimal payments are decreasing with time. We conclude with some discussion of practical considerations for real-world implementation.

This work contributes to a growing literature on payments for environmental goods [Engel *et al.*, 2008; Jayachandran *et al.*, 2017; Fenichel *et al.*, 2019]. We also identify a set of future directions in the space of market design for environmental sustainability for AI and economics community.

Outline. We set up the model in Section 2, describe the agent's strategy in Section 3 and solve the principal's problem in Section 4. We conclude by discussing limitations and extensions in Section 5.

2 Model

State space. We use a Markov chain with finite state space $\mathcal{S} = \{0, 1, \dots, M\}$ to model the state of a tree. M is the number of periods (years) that a tree takes to mature. State 0 represents no tree, states 1 to $M - 1$ represents the growing process and the final state M represents maturity. At time step $t = 0$, the agent starts in state $s = 0$. We assume $M \geq 2$, so that there is at least one intermediate state before reaching maturity. We assume the principal has a monitoring technology available, so the state is publicly observable.

Agent actions. In every period, the agent chooses a binary action a from set $\mathcal{A} = \{0, 1\}$ where $a = 0$ means no effort and $a = 1$ means exerting effort. The action is *unobservable* to the principal. Action $a = 0$ is costless but action $a = 1$ costs $c_s, \forall s \in \{0, \dots, M - 1\}$ and 0 at maturity state M .¹ For most of the paper we will assume *constant cost* per state; that is, each agent incurs cost c for every single state. The agent type c is drawn i.i.d. from CDF F with support $[0, \bar{c}]$. The constant cost assumption is for technical convenience; we will later demonstrate that our main results apply more generally. The principal has the knowledge of the distribution $F(c)$ but not the agent's individual type.

Exogenous shocks. In every period, a tree might die due to some natural risk out of the farmer's control. We model this risk as a probabilistic exogenous shock- in each period where the tree is in state s , a shock does *not* occur with probability

¹Note that we could also interpret $a = 0$ as actively cutting down the tree to use or sell, in which case c_s captures the outside option cost of keeping the tree.

q_s . For most of the paper we assume constant shock probability $q_s = q$, but our results extend to state-dependent shocks (such as younger trees being more vulnerable to natural risk).

Transition probability. The transition probabilities depend on the agent's behavior and exogenous shocks. If, at time t , the state is $s_t \in \{0, 1, \dots, M\}$, then $s_{t+1} = \min\{s + 1, M\}$ (i.e., the tree grows or stays at maturity if already mature) if and only if the agent chooses $a = 1$ and no exogenous shock occurs. Otherwise, $s_{t+1} = 0$ (i.e., the tree is lost) if $a = 0$ or if the shock occurs.

Value at maturity. In state s , the principal and the agent receives values v_s^P and v_s^A , respectively, from a tree. We assume values for both parties from an immature tree are zero - $v_s^P = 0$ and $v_s^A = 0, \forall s \in \{0, \dots, M - 1\}$. A mature tree delivers values $v_M^P \geq 0$ and $v_M^A \geq 0$.

Principal's payment. In each round the principal can transfer a payment to the agent. Payments from the principal can depend on the state transition but not the agent's action. We without loss restrict attention to payments $\{p_s\}_{s \in \mathcal{S}} \in \mathbb{R}_+^{M+1}$, where each p_s is a payment conditional on a forward transition of reaching state $s + 1$ and the payment conditional on transitioning to state 0 is 0; p_M is the payment the agent receives for keeping mature trees alive.

Agent utility. The agent is risk neutral and his utility is linear in the payment. In what follows, we assume the action choice of the agent is Markovian, i.e., does not depend on the history. This assumption is without loss in Markovian games. In state s , if the agent chooses action $a = 0$, his current state utility $u_s(a = 0)$ is 0. If he chooses action $a = 1$, his stage utility $u_s(1)$ is $p_s + v_s^A - c_s$ if the tree survives and $-c_s$ otherwise; together, his expected current state utility is $u_s(a = 1) = q_s(p_s + v_s^A) - c_s$.

The agent discounts the future with discounting rate $\delta < 1$. Assuming the agent chooses action a_s in state s , his continuation utility is $\sum_{t=1}^{\infty} \delta^t u_{\sigma(s,t)}(a_{\sigma(s,t)})$ where $\sigma(s, t)$ is a random variable denoting the state after t transitions from state s , giving a total expected utility from state s of:

$$\mathbb{E}[u_s(a_s) + \sum_{t=1}^{\infty} \delta^t u_{\sigma(s,t)}(a_{\sigma(s,t)})]. \quad (1)$$

We assume that the agent always chooses the option preferred by the principal when the agent is indifferent.

2.1 The Principal's Problem

The principal's objective is to maximize the long-run expected value from mature trees minus the expected cost of incentive payments:

$$\max_{\{p_s\}} \mathbb{E}_{\sigma \sim \mathcal{D}} [v_{\sigma} - p_{\sigma}] \quad (2)$$

where the expectation is over the steady state distribution \mathcal{D} of the Markov chain. Our focus on the steady-state is motivated by the fact that, in practice, the principal will be interacting with many agents who have many trees, in which case the steady-state of the process is a proxy for the aggregate outcome. Given a choice of $\{p_s\}$, the principal's payoff

is determined by the agent's utility-maximizing choice of actions, which is endogenous to this payment and induces the steady state of the Markov chain. Last, we assume the principal faces a limited liability constraint, which in this setting means that all payments are non-negative.

3 The Agent's Perspective

The agent can choose between two actions in every period, but we note that due to the structure of the Markov process it suffices to consider the following set of stationary strategies.

Definition 1 (The agent strategy). *The set of strategies, denoted as $\phi \in \{0, \dots, M, \infty\}$, correspond to choosing to exert effort only up to a certain state. Explicitly, in strategy $\phi \in \{1, \dots, M\}$ the agent chooses to exert effort ($a = 1$) in states $s \in \{0, \dots, \phi - 1\}$ and otherwise $a = 0$. Strategy $\phi = 0$ corresponds to choosing $a = 0$ for all states (not participating); and $\phi = \infty$ corresponds to choosing $a = 1$ in every state.*

For each strategy ϕ , the expected cost, evaluated at some state s , is the sum of the current period cost and the discounted expected future cost of choosing ϕ , denoted as $\mathbb{E}C_s^\phi(c)$. We can compute this cost for each $\phi < \infty$ by solving a set of $\phi + 1$ linear equations following the Markov chain dynamics. We have

$$\mathbb{E}C_s^\phi(c) = c + \delta q \mathbb{E}C_{s+1}^\phi(c) + \delta(1 - q) \mathbb{E}C_0^\phi(c), \forall s \in \{0, \dots, \phi - 1\} \quad (3)$$

$$\mathbb{E}C_\phi^\phi(c) = \delta \mathbb{E}C_0^\phi(c) \quad (4)$$

$$\Rightarrow \mathbb{E}C_0^\phi(c) = K_\phi c, \forall \phi \in \{1, \dots, M\}, \quad (5)$$

where the coefficients K_ϕ depend on q, δ, ϕ (exact expression omitted for space reason). $\mathbb{E}C_0^\infty(c)$ can be computed by adjusting Eq. (4) in the above linear equations to allow the transition from state M to itself. We include $K_0 \equiv 0$.

For notational convenience, we define $\mathbb{E}C^\phi(c) \equiv \mathbb{E}C_0^\phi(c)$ to be the expected total (future discounted) cost for the agent c in equilibrium, as evaluated at state 0. In this constant cost model, we are able to observe:

$$K_0 < K_1 < \dots < K_M, \quad K_M > K_\infty \quad (6)$$

In other words, strategies $\phi \leq M$ get more costly as ϕ increases, since the agent spends more time exerting effort in expectation. The exception is $\phi = \infty$, which is less costly than strategy $\phi = M$ because the agent can keep the tree mature without any cost.

Similarly, we can calculate the agent's expected total payments from a given payment plan and from having mature trees. Given any payment plan, the expected payments of strategy ϕ is denoted $\mathbb{E}B^\phi(p_0, \dots, p_{\phi-1})$. Note that $\mathbb{E}B^0 = 0$ because no payment is received when not enrolling.

We denote the expected value of a tree when the agent chooses strategy ϕ to be $\mathbb{E}v_A^\phi$, where $\mathbb{E}v_A^\phi = 0$ for all strategies except for $\phi = \infty$. $\mathbb{E}v_A^\infty$ is proportional to v_M^A with some coefficients based on the Markov chain.

Given a payment plan $\{p_s\}$, the agent chooses a strategy $\phi \in \{0, \dots, M, \infty\}$ that maximizes his expected utility: $\mathbb{E}U^\phi = \mathbb{E}B^\phi(\{p_s\}) - \mathbb{E}C^\phi(c) + \mathbb{E}v_A^\phi$.

4 Solving The Principal's Problem

We start off by formulating the constraints faced by the principal. Next, we investigate a reduced problem: find the least-costly contract such that an agent with cost $c \in [c_l, c_h]$ chooses $\phi = \infty$ (and whether any such contract exists). We show how to compute the optimal payment schedule for this reduced problem. We then reduce the more general original problem in Eq. (2) to this reduced problem, by finding an optimal subinterval of the agent type to target and solving the reduced problem for that subinterval.

In order for the agent with cost c to choose $\phi = \infty$, we require the following M incentive compatibility constraints (IC): $\forall \phi \in \{1, \dots, M\}$,

$$\mathbb{E}B^\phi(\{p_s\}) - \mathbb{E}C^\phi(c) \leq \mathbb{E}B^\infty(\{p_s\}) - \mathbb{E}C^\infty(c) + \mathbb{E}v_A^\infty. \quad (7)$$

We also require the individual rationality constraint (IR) for agent c choosing strategy $\phi = \infty$:

$$\mathbb{E}B^\infty - \mathbb{E}C^\infty(c) + \mathbb{E}v_A^\infty \geq 0. \quad (8)$$

4.1 Who Drops Out and When

One might expect that if a contract causes the agent with a high cost to follow through (i.e., choose strategy ∞), one with a lower cost would follow through as well. However, this is not always the case: as we next show, the agent with lower costs might be incentivized to drop out of the program "late" (i.e., just before the tree reaches maturity) then reenter with a new tree in order to collect more payments.

We now make this more formal. Denote $\mathbb{E}\bar{B}^\phi(c)$ to be the minimal expected payments that type c needs to receive in order to weakly prefer strategy ϕ to ∞ . By definition, $\mathbb{E}\bar{B}^\phi(c) \equiv \mathbb{E}C^\phi(c) - \mathbb{E}C^\infty(c) + \mathbb{E}B^\infty + \mathbb{E}v_A^\infty$. The term $\mathbb{E}\bar{B}^\phi(c)$ measures how much expected payments an agent requires in order to cause deviation from the principal's desired strategy, ∞ . The following lemma demonstrates that high-cost types are prone to choosing early drop-out strategies and low-cost types are prone to choosing late drop-out strategies.

Lemma 1. *For any pair of $c_l < c_h$ in the support of $F(c)$, there exists a state $\hat{s} \in \{0, \dots, M\}$ such that $\forall \phi < \hat{s}, c_l = \arg \min_{c \in [c_l, c_h]} \mathbb{E}\bar{B}^\phi(c)$ and $\forall \phi \geq \hat{s}, c_h = \arg \min_{c \in [c_l, c_h]} \mathbb{E}\bar{B}^\phi(c)$.*

Proof sketch. We show that $\mathbb{E}\bar{B}^0(c_l) > \mathbb{E}\bar{B}^0(c_h)$ and $\mathbb{E}\bar{B}^M(c_l) > \mathbb{E}\bar{B}^M(c_h)$. Further, we know that $\mathbb{E}\bar{B}^\phi(c)$ increases in ϕ . There must be an intermediate strategy where the type that requires the minimal expected benefits switches from high to low. The full proof is omitted due to space constraints.

4.2 How Much to Pay

We now construct a contract that minimizes the expected total payments required to have a given subpopulation always complete the tree-growing process. Our first observation is that if the agent with costs c_l and c_h choose $\phi = \infty$, then any agent with cost $c \in [c_l, c_h]$ will choose it as well. It therefore suffices to consider subpopulations corresponding to cost intervals. Theorem 1 describes the optimal payments for a given cost interval.

Theorem 1. [Optimal payment schedule] Given any subpopulation $[c_l, c_h] \subseteq \text{supp}(F)$, the least cost contract is the payment schedule, p_0^*, \dots, p_M^* , that solves the following set of $M + 1$ equations. For all $s \in \{1, \dots, \hat{s} - 1\}$:

$$\mathbb{E}B^s(p_0, \dots, p_{s-1}) = \mathbb{E}C^s(c_h), \quad (9)$$

Otherwise:

$$\mathbb{E}B^s(p_0, \dots, p_{s-1}) = \mathbb{E}C^s(c_l) + \mathbb{E}B^\infty + \mathbb{E}v_A^\infty - \mathbb{E}C^\infty(c_l), \quad (10)$$

And,

$$\mathbb{E}C^\infty(c_h) = \mathbb{E}B^\infty + \mathbb{E}v_A^\infty. \quad (11)$$

Proof intuition. The proof is composed of two parts. First, the principal always improves the objective by switching payments to earlier states as long as the constraints in Eq. (7) are satisfied. This is due to the difference in discounting factors between the agent and the principal. Second, we identify a small set of binding constraints. We reduce the number of constraints from $(|c| + 1) \times (M + 1)$ to $M + 1$ by utilizing the analysis in Section 3 and Lemma 1.² The full proof is omitted due to space constraints.

In the rest of this section, we discuss the implications of the optimal payment schedule described in Theorem 1.

Binding constraints. First, we discuss which constraints are binding in the optimal solution. It’s intuitive to see that in the optimal solution, c_h ’s IR constraint binds (Eq. (11)). The agent c_h is indifferent between $\phi = \infty$ and not participating. It’s sufficient to have IR satisfied for c_h because IR will also be satisfied for all lower types. Regarding IC constraints, although the principal prefers to shift payments to earlier states, there is a limit on how much shifting is possible. As we increase early stage payments ($p_s, \forall s \in \{0, \dots, \hat{s} - 2\}$), the IC constraint of the agent c_h will be violated first; this causes c_h to choose early drop-out strategies (Eq. (9)). Similarly, late-stage payments ($p_s, \forall s \in \{\hat{s} - 1, \dots, M\}$) cannot exceed the minimal payments that keeps low-type c_l indifferent between dropping out in intermediate stages and strategy $\phi = \infty$ (Eq. (10)). The threshold state \hat{s} comes from Lemma 1.

Take-up and follow-through behavior. Given the optimal payment schedule, all intermediate types in $[c_l, c_h]$ choose $\phi = \infty$, any types higher than c_h choose ϕ_0 , and the types lower than c_l choose ϕ_M . Even though the principal only intends to have $c \in [c_l, c_h]$ to take up and follow through the contract, she cannot prevent lower types $c < c_l$ from taking up but not following through. This is consistent with the self-selection observations made in [Jack and Jayachandran, 2018]. Further, the low cost agent ($c < c_l$) chooses to drop out before the tree reaches maturity even in the absence of exogenous shocks. If the contract is not properly designed, then this intentional drop out behavior will be exacerbated. In [Jack *et al.*, 2015], the authors argue that exogenous shocks cause participants to not follow through, thus lowering program cost-efficiency. While this is consistent with our

² $|c|$ can be thought as the size of the discretization of the agent type space between c_l and c_h ; there are M IC constraints and 1 IR constraint for each type; there are $M + 1$ number of limited liability constraints

model, we further contribute to this discussion by showing that another possible reason for drop-out is that front-loaded payments (even the optimal payments) can incentivize some agent type to join the program to collect early payouts but then abandon trees.

Finding the optimal interval. To solve the optimization problem with objective in Eq. (2), we can simply search for the optimal interval $[c_l, c_h]$ of the agent type that will choose $\phi = \infty$, where the expected payment is given by Theorem 1. This can be done efficiently via discretization plus grid search. We omit details due to space constraints.

Payment to keep a mature tree. In the optimal payment schedule, the final state payment p_M rewards transitions from M to M . In an afforestation program, it’s intuitive to compensate a participant for tree growing efforts which involve non-zero payments in p_0, \dots, p_{M-1} . Once a tree matures, it is no longer costly for the agent to keep the tree alive. A positive payment p_M may appear to be unnecessary but in the optimal solution, the principal may have to keep paying the agent even after a tree reaches M so that the agent does not cut it down and reenter the program.³ However, the principal does not need to pay the agent a positive price to keep a mature tree if the value of a mature tree to the agent is large.

Corollary 1. For every c_h , there is a v_M^A that is large enough such that the optimal payment schedule will induce $\forall c \leq c_h$, c to choose $\phi = \infty$ and $p_M^* = 0$.

Although v_M^A is exogenous in our model, in reality it can be partially affected by design. By working with local stakeholders, the principal can offer tree seedling options preferred by farmers and educate them the benefit of agroforestry (for example, the “Trees on Farms for Biodiversity” project at the World Agroforestry Center⁴).

5 Limitations & Extensions

In our model, the agent’s cost type is the sole source of heterogeneity. Other potential sources of heterogeneity to consider include the agent’s discounting rate and risk preference [Ihli *et al.*, 2016; Jack, 2013]. We assume that the principal has the knowledge of the agent cost distribution throughout the tree growing stages. This further advocates the need for robustness considerations.

We also consider a model with stochastic agent costs, where in every period, the agent cost is redrawn from some distribution. The stochastic agent cost model aims to answer how payments can be designed to alleviate dropout due to stochastic shocks, including income shocks.

Acknowledgments

The authors want to thank Microsoft AI for Earth and the World Agroforestry Center (ICRAF). This research is partly funded by NSF EAGER CCF-1841550: Algorithmic Approaches for Developing Markets.

³In reality, a low-cost agent’s reentering decision may look like planting more trees. In our model, we consider an agent who have reached his land capacity. Thus, the agent only reenter through cutting down existing trees and planting new ones.

⁴<https://treesonfarmsforbiodiversity.com/about-trees-on-farms/>

References

- [Carroll, 2015] Gabriel Carroll. Robustness and linear contracts. *Am. Econ. Rev.*, 105(2):536–563, 2015.
- [Dao *et al.*, 2019] David Dao, Catherine Cang, Clement Fung, Ming Zhang, Reuven Gonzales, Nick Beglinger, and Ce Zhang. GainForest : Scaling Climate Finance for Forest Conservation using Interpretable Machine Learning on Satellite Imagery. *ICML Clim. Chang. Work. 36th Int. Conf. Mach. Learn.*, pages 5–7, 2019.
- [DeMarzo and Sannikov, 2006] Peter M. DeMarzo and Yuliy Sannikov. Optimal security design and dynamic capital structure in a continuous-time agency model. *J. Finance*, 61(6):2681–2724, 2006.
- [Dütting *et al.*, 2019] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. Simple versus optimal contracts. *ACM EC 2019 - Proc. 2019 ACM Conf. Econ. Comput.*, pages 369–387, 2019.
- [Engel *et al.*, 2008] Stefanie Engel, Stefano Pagiola, and Sven Wunder. Designing payments for environmental services in theory and practice: An overview of the issues. *Ecol. Econ.*, 65(4):663–674, 2008.
- [Fenichel *et al.*, 2019] Eli P Fenichel, Wiktor Adamowicz, Mark S Ashton, and Jefferson S Hall. Incentive Systems for Forest-Based Ecosystem Services with Missing Financial Service Markets. *J. Assoc. Environ. Resour. Econ.*, 6(2):319–347, 2019.
- [Hansen *et al.*, 2013] M. C. Hansen, A. Tyukavina P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, and J. R. G. Townshend D. Thau, S. V. Stehman, S. J. Goetz, T. R. Loveland, A. Kommareddy, A. Egorov, L. Chini, C. O. Justice. High-Resolution Global Maps of 21st-Century Forest Cover Change. *Science (80-.)*, 342(November):850–854, 2013.
- [Holmstrom and Milgrom, 1987] Bengt Holmstrom and Paul Milgrom. Aggregation and Linearity in the Provision of Intertemporal Incentives. *Econometrica*, 55(2):303–328, 1987.
- [Ihli *et al.*, 2016] Hanna Julia Ihli, Brian Chiputwa, and Oliver Musshoff. Do changing probabilities or payoffs in lottery-choice experiments affect risk preference outcomes? Evidence from rural Uganda. *J. Agric. Resour. Econ.*, 41(2):324–345, 2016.
- [IPCC, 2019] IPCC. Climate Change and Land: an IPCC special report on climate change, desertification, land degradation, sustainable land management, food security, and greenhouse gas fluxes in terrestrial ecosystems, 2019.
- [Jack and Jayachandran, 2018] B. Kelsey Jack and Seema Jayachandran. Self-selection into payments for ecosystem services programs. *Proc. Natl. Acad. Sci.*, page 201802868, 2018.
- [Jack *et al.*, 2015] B. Kelsey Jack, Paulina Oliva, Christopher Severen, Elizabeth Walker, and Samuel Bel. Technology adoption under uncertainty: Take up and subsequent investment in Zambia. *Natl. Bur. Econ. Res.*, 21:52, 2015.
- [Jack, 2013] Kelsey K. Jack. Private information and the allocation of land use subsidies in Malawi. *Am. Econ. J. Appl. Econ.*, 5(3):113–135, 2013.
- [Jayachandran *et al.*, 2017] Seema Jayachandran, Joost De Laat, Eric F Lambin, Charlotte Y Stanton, Robin Audy, and Nancy E Thomas. Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation. *Science (80-.)*, 357(6348):267–273, 2017.
- [Jean *et al.*, 2019] Neal Jean, Sherrie Wang, Anshul Samar, George Azzari, David Lobell, and Stefano Ermon. Tile2Vec: Unsupervised Representation Learning for Spatially Distributed Data. *Proc. AAAI Conf. Artif. Intell.*, 33:3967–3974, 2019.
- [Levin, 2003] Jonathan Levin. Relational Incentive Contracts. *Am. Econ. Rev.*, 93(3):835–857, 2003.
- [Lütjens *et al.*, 2019] Björn Lütjens, Lucas Liebenwein, and Katharina Kramer. Machine Learning-based Estimation of Forest Carbon Stocks to increase Transparency of Forest Preservation Efforts. *NeurIPS Work. Tackling Clim. Chang. with AI*, 2019.
- [Ross, 1973] S.A. Ross. The Economic Theory of Agency : The Principal ’ s Problem Linked references are available on JSTOR for this article : The Economic Theory of Agency : The Principal ’ s Problem. *Am. Econ. Rev.*, 63(2):134–139, 1973.