AgBioForum, 13(1): 13-24. ©2010 AgBioForum.

Stochastic Benefit Streams, Learning, and Technology Diffusion: Why Drought Tolerance is Not the New Bt

Travis J. Lybbert and Adrian Bell
University of California, Davis

The speed of Bt cotton diffusion among smallholders in poor countries such as India, China, and South Africa has been unprecedented. Hopes are high for drought-tolerant (DT) varieties that similarly reduce yield risk and have attracted substantial investments from public, private, and philanthropic sectors. We highlight important learning differences between Bt and DT that will shape diffusion patterns. While the potential welfare benefits of DT are compelling, we caution against glossing over practical complications that farmers will face in assessing the relative merits of DT varieties. We emphasize how and why vulnerable farmers facing marginal growing conditions—ostensibly, the target beneficiaries—may be slowest to adopt DT crops. More frequent extreme drought events associated with climate change may further complicate this learning and diffusion process. Generalized water-use efficiency gains and early maturation could help improve learning in arid and semi-arid regions by conferring benefits across a broader range of rainfall outcomes.

Key words: drought tolerance, adoption, learning, risk aversion, climate change, Bt cotton.

Introduction

The Green Revolution brought spectacular yield gains to many crops in many parts of the developing world. In the current generation of agricultural researchers, the potential of agricultural biotechnology to produce seed traits that reduce risk—rather than unconditionally increasing yields—generates real excitement. Far and away the most successful such technology involves pest resistance conferred by the *Bacillus thuringiensis* (Bt) gene. The speed of diffusion of this technology has been unprecedented. The benefits to smallholders in India appear to be substantial (e.g., Qaim & Zilberman 2003; Subramanian & Qaim, 2009). Although the factors influencing adoption differ across countries (Smale et al., 2009; Tripp, 2009), nearly every cotton farmer in some regions has adopted Bt cotton less than a decade after it was introduced. In the wake of this roaring success, hopes are high for drought-tolerant (DT) varieties that similarly reduce yield risk.

Few agricultural research objectives have ever attracted the intensity of attention and investment from private, public, academic, and philanthropic sectors as DT has in recent years. All the major private and public players in agricultural research have invested in DT, as have major philanthropies such as the Bill and Melinda Gates Foundation. In the past decade, total investments in DT research almost certainly have surpassed $1 billion. This level of interest in DT appears well justified: with climate change, growing water scarcity and impending water disputes, the prospective welfare gains from DT varieties are enormous. In the United States and Australia, these gains translate into significant market potential. In poor countries, gains from DT could mean the difference between survival and starvation. During a drought, DT may protect vulnerable households from catastrophic losses and help them recover more quickly from such a shock. Many proponents point to an additional, more subtle entrepreneurial spillover benefit. Since DT reduces production risk, poor households that adopt DT varieties may feel less vulnerable and may therefore shift other parts of its livelihood portfolio into slightly higher risk, but higher return, pursuits. If it occurs in practice the way proponents claim it might, this "positive moral hazard" would further magnify the potential welfare gains to DT for households with few other risk-management options.

Impressive as these potential DT benefits are, critical adoption considerations seem to be ignored or dismissed by an implicit 'build it and they will come' mentality. This perennial pitfall in disseminating new technologies in developing countries may be less appreciated in the wake of the unprecedented success of Bt crops among the poor, but it is no less relevant. We highlight DT adoption constraints by contrasting Bt and DT crops. We capture India's Bt cotton experience in a stylized learning and adoption model. We modify the model to reflect DT characteristics and showcase key differences between Bt and DT. We then explore how farmer vulnerability, marginality of growing conditions, seed costs, and climate change may further slow the diffusion of DT. While the potential welfare benefits of DT are huge, the learning hurdles that complicate DT adoption and diffusion deserve more careful attention.

Stochastic Benefits and Learning

While the diffusion of new seed varieties is often shaped by institutions, culture, and market structure, the underlying adoption decision is fundamentally a learning process. We acknowledge the importance of institutional and other determinants of diffusion (see Smale et al., 2009; Tripp, 2009), but restrict our focus here to learning and adoption. Learning may begin with purported benefits claimed on a label or by an agro-services dealer or extension agent. The most influential learning, however, involves personal experimentation and learning from others. Learning from experience—whether one's own or that of an acquaintance—is nowhere more important than with the highly stochastic and heterogeneous agricultural production of smallholder agriculture in developing countries (Conley & Udry, 2010; Foster & Rosenzweig, 1995; Marra, Pannell, & Abadi, 2003). In these settings, the adoption decision typically involves marginal adjustments in response to observed benefits of a new variety relative to some benchmark variety. While this approach often loosely follows an implicit treatment-control design, fluctuations and constraints out of the farmer's control produce background noise in this experimentation process. This noise can be deafening for poor farmers on marginal land who face greater fluctuations and constraints—e.g., due to poor soil quality and limited access to irrigation, fertilizer, credit, and risk-management options.

Given these learning challenges, widespread diffusion can only occur when new varieties confer a strong enough relative benefit to stand out amidst this background noise. In many places, the high-yielding hybrids of the Green Revolution conferred such a strong signal. These hybrids had three features that facilitated this learning process: (i) they marked a discrete and dramatic improvement over benchmark varieties, (ii) they conferred a nearly static relative benefit stream that was not strongly conditioned on underlying environmental stochasticity, and (iii) they often aimed at irrigated settings with relatively less background noise. These features imply that every harvest provided meaningful information with which to update one's valuation of the new seeds. This enables farmers to quickly learn what the new technology is worth to them. With limited background noise, learning from others relatively straightforward, and this social learning creates important diffusion dynamics that can magnify the effect on adoption patterns of a dramatic, nearly unconditional relative benefit stream.

Seed traits that reduce risk are both popular and promising, but they also complicate the learning process—particularly for the poor farmers in developing countries who are often the target clientele. The benefit streams of these traits involve changes in higher moments of the yield distribution, which make their relative benefit streams stochastic. The information content of a given harvest is thus conditioned on the underlying target stress (e.g., pests, drought). In some seasons, the relative merits of the new seed are obvious, but in others there may be no difference at all. These stochastic relative benefit streams complicate rigorous scientific experimentation and make documenting the merits of DT traits challenging even for plant breeders exerting significant control over growing conditions. Add substantial background noise to this learning process, and the experimentation task of the smallholder looks daunting indeed. Stochastic relative benefit streams may make diffusion more erratic with swings of adoption and disadoption as farmers update their assessment of the new technology based on previous season experiences. Here the quasi-insurance function of risk reducing seeds is relevant. Bt and DT seeds may indeed function like insurance with a payout some years and a lost 'premium' payment others, but agronomic heterogeneity and environmental stochasticity effectively conceal the terms of this insurance from the farmer. Throughout the learning process, these concealed insurance terms mute any spillover *ex-ante* benefits of reduced risk exposure. Discovering the terms of this built-in insurance is precisely the process of learning the relative merits of risk-reducing seeds.

Bt Cotton in India and How DT Differs

There are important differences between the stochastic relative benefit streams conferred by different risk-reducing traits. We characterize Bt cotton in India and then compare it to DT in order to highlight learning complications that are unique to DT. Few agricultural technologies have been tracked in their diffusion as closely and with as much interest as Bt cotton has in India (e.g., Qaim, 2003; Qaim, Subramanian, Naik, & Zilberman, 2006; Qaim & Zilberman, 2003; Smale et al., 2009; Subramanian & Qaim, 2009; Tripp, 2009). The story, which has been documented and spun from a variety of perspectives (some more biased than others [Herring, 2008]), demonstrates some of the learning complications we model. Rather than duplicating efforts to describe this experience in detail, we highlight a few features that are particularly relevant for learning and adoption.

In 2001, cotton growers in the state of Gujarat were ravaged by bollworms—with the exception of those growing illegal (i.e., unapproved) Bt cotton hybrids. The discovery and demonstration effect of Bt cotton sparked an explosion of interest from farmers and forced India's Genetic Engineering Approval Committee to act fast to approve the Mahyco (Monsanto licensee) Bt cotton hybrids that had already undergone official testing for several years. While 2002 marked the official launch of Bt cotton in India, the market was awash in illegal Bt cotton seeds (both real and fake). Mahyco priced its approved Bt hybrids at a 300% premium over its non-Bt hybrids, which further fueled the black market. Bollworm pressure was much lower that year. Many farmers felt the Bt seeds were not worth the premium they had paid. This led some to disadopt in 2003, but diffusion quickly recovered as farmers learned more about the stochastic relative benefits of Bt cotton. Since then, most years have seen rapid expansions in diffusion and in the range of approved Bt cotton hybrids available on the market. In 2006, the combination of price controls and a flood of new Bt hybrids entering the market squelched the black market for Bt cotton and lowered the average Bt cotton premium to less than 90%—although this premium is becoming less meaningful because many dealers are not even stocking non-Bt cotton seeds anymore.

What explains this remarkably rapid learning and diffusion process? Bt crops have built-in pest protection. A Bt cotton hybrid differs from a non-Bt version of the same hybrid only in its production of a toxin that kills bollworms that eat any part of the plant. The relative benefit stream conferred by Bt cotton is therefore a function of pest pressure. Since a tiny amount of green tissue ingestion is sufficient to kill a bollworm, the protection offered by Bt is effective even under extreme pest pressure. In the short run, bollworms simply cannot overpower a Bt cotton plant.[1] This has an important implication for learning: the relative benefit stream and hence information content of Bt cotton is monotonically increasing in pest pressure.

1. *This ignores the possibility of resistance build up—a theoretical possibility that has yet to be documented empirically.*

In many developing countries, controlling bollworms via pesticides is imprecise and imperfect. Consequently, nearly every season brings some degree of pest damage. The comparison with the complete protection conferred by Bt can be stark. When baseline levels of uncontrolled pest damage are high in such settings, Bt cotton can therefore confer an unconditional benefit in addition to the stochastic one. For example, baseline bollworm control in South Africa is quite imperfect, and Bt cotton seems to confer a significant unconditional yield advantage (Qaim & Zilberman, 2003; Thirtle, Beyers, Ismael, & Piesse, 2003). In China, however, baseline control was much better, and Bt cotton conferred little unconditional yield advantage (Pray, Ma, Huang, & Qiao, 2001). India seems to be much more like South Africa in this regard, and most Indian farmers have seen an obvious advantage in most years.

Table 1 displays key differences between Bt and DT that are relevant to learning and subsequent adoption in developing countries. Although upstream differences in complexity in Rows A and B may matter most to scientists, they can also complicate farmers' learning processes by making relative benefits sensitive to other interactions and timing effects. Once seeds with these traits arrive on farmers' fields, other key differences emerge. We highlight eight such differences in Table 1, a few of which merit some discussion. First, drought stress directly affects households in many ways, whereas bollworm pressure may only directly affect their cotton crops—even if this narrow direct effect decimates the crop. This means that the relative benefits of DT under extreme drought stress matter more than do the benefits of Bt under a comparably high degree of pest pressure. In conjunction with the second difference, this has important implications for learning as we demonstrate in our model. Next, baseline bollworm control is low in many developing countries, which implies that a portion of the yield benefits of Bt is nearly unconditional (3). DT is much less likely to confer unconditional yield gains and may even suffer a yield penalty in the absence of drought (4 and B). Effective pesticide treatments can reduce yield differences between Bt and non-Bt crops in a given year (6). Substantial performance differences can still appear in net returns, but these returns are much harder to observe than yield differences. In rainfed agriculture, there is no comparably effective remedy to drought stress. While in isolation this may make it easier for farmers to learn about DT than Bt, substantial temporal and spatial background noise make it much more difficult to attribute observ-

Table 1. Differences between Bt and DT that are relevant for farmer learning.

			Bt	Drought tolerance
'Upstream' differences	A	Effect of target stress on plant physiology	Simple	Complex: "Drought stress is as complicated and difficult to plant biology as cancer is to mammalian biology" (Jian-Kang Zhu, as cited by Pennisi, 2008, pp. 171). Timing of stress is crucial (e.g., stress at flowering can mean zero yield).
	B	Transmission of trait	Single gene to "flip switch"	Notoriously difficult with a complex network of metabolic and physiological pathways that require scientists to "dial into the physiology of the plant" (Pennisi, 2008, pp. 171).
'Downstream' differences	1	Sensitivity of general household welfare to target stress	Low	High
	2	Stochastic relative benefit stream	Monotonically increasing in pest pressure: Maximum pressure is best for learning about Bt	Non-monotonic in drought pressure: Moderate drought is best for learning about DT. Severe drought makes DT and non-DT indistinguishable.
	3	Unconditional yield gain	Potentially high depending on efficacy of farmers' baseline bollworm control	Low to zero
	4	Temporal background noise: Cumulative effects of stress	Low: Cyclical but not cumulative	High and strongly conditioned on soil moisture, soil quality, organic matter, plot slope, rotation history, access to supplementary irrigation, etc.
	5	Spatial background noise: Degree of spatial heterogeneity in stress	Low: Pest pressure within a village and even at larger spatial scales relatively homogeneous	High: Microclimates and soil heterogeneity can make it difficult to discern subtle differences across plots even within a village.
	6	Performance penalty of trait	Zero (see B above)	Uncertain but potentially high (see B above)
	7	Purchase price of trait	High initially, but falling rapidly	Low to zero in poor countries due to "no royalty" agreements (but effective seed delivery may be difficult nonetheless).
	8	Pre-trait remedy to target stress	Simple and direct: Pesticide application	Complex and indirect: With few options for directly remedying drought stress, rainfed farmers resort to adjusting production strategies and asset portfolios.

able yield differences to DT as opposed to many other potentially confounding factors (7 and 8).

Learning and Diffusion Model

There are important differences between the stochastic relative benefit streams conferred by different risk-reducing traits. A comparison of Bt and DT reveals several such differences, many of which imply differences in learning, adoption, and diffusion. We use a few of these Bt-DT differences to pin down a learning and diffusion model. We then focus the model on DT to further explore implications of these learning complications.

Relative Yield Gains and Utility Benefits

Suppose that consumption and production are separable so that utility can be written as the constant relative risk-aversion function $u(z,w) = \dfrac{(z+w)^{1-\alpha}}{1-\alpha}$, where α is the coefficient of relative risk aversion and income is composed of the yield z of the target crop and other income w. We normalize the output price of the crop so that both z and w are measured in monetary units. The target stress is measured as an index $x \in [0, 1]$ such that baseline yield is $z(x) = \bar{z}(1 - x)$. The yield of the new varieties as a function of x reflects the fundamental differences highlighted in Table 1. Individuals choose between a baseline variety with this yield response function and a new variety (first Bt, then DT) that is less sensitive to x. The Bt variety is impervious to target pest pressure, so its yield response function is $z_{Bt}(x) = \bar{z}_{Bt}$. While $\bar{z}_{Bt} > \bar{z}$ reflects an unconditional yield benefit due to incomplete baseline pest management (Qaim & Zilberman, 2003), net benefits will not be unconditionally higher with Bt if farmers must pay a premium for the Bt variety. In contrast, the yield response function of the DT variety is non-monotonic in drought pressure:

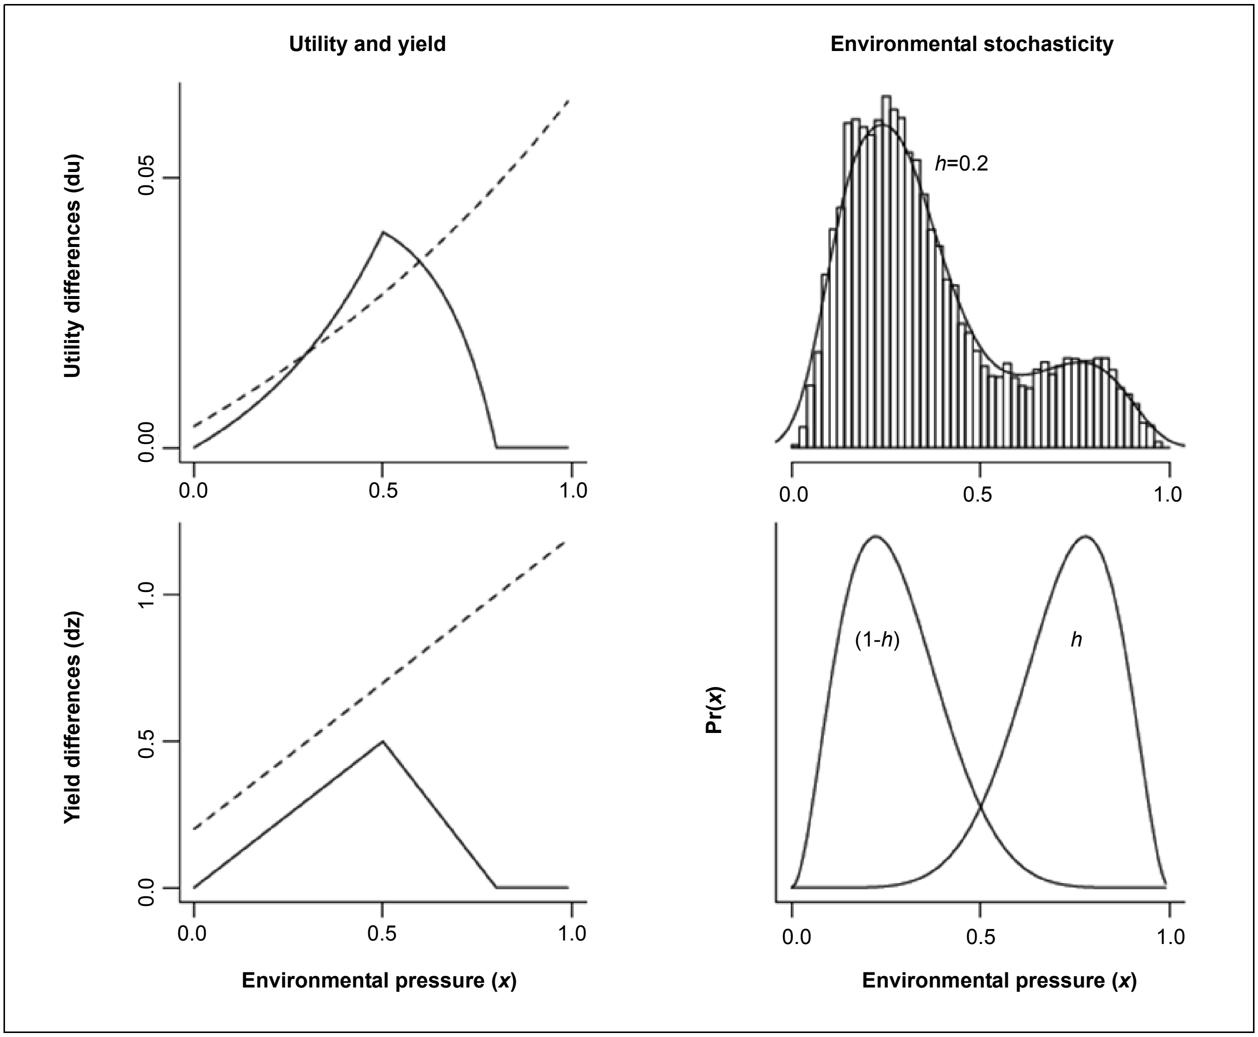

Figure 1. Relative yield gain and relative utility benefit by trait (DT solid, Bt dashed; left panels) and probability distribution of environmental pressure (right panels).

$$z_{DT}(x) = \begin{cases} \bar{z} & if \quad x \in [0, x_1] \\ \bar{z} - g(x) & if \quad x \in [x_1, x_2] \\ \bar{z}(1-x) & if \quad x \in [x_2, 1] \end{cases} \tag{1}$$

In contrast to Bt, a DT variety can only tolerate drought stress to a point (x_1) before it too suffers drought-induced yield losses ($g(x)$). At extreme drought pressure (x_2), the DT variety is indistinguishable on average from its non-DT counterpart. Bt seeds cost $c>0$ more than the baseline variety, but DT seeds have no such premium (Table 1). To capture the perspective of a farm household, we must account for how the underlying target stress affects household welfare beyond the yield response functions above. We account for this in other household income and assume that $w(x) = \bar{w}(1 - \delta_i x)$, where δ_i indicates the sensitivity of other household income to stress, i={Bt, DT}, and $\delta_{Bt} < \delta_{DT}$.

In season t, the relative net yield gain due to Bt is $dz_{Bt}\,(x_t) = \bar{z}_{Bt} - z(x_t) - c$. The analogous relative net yield gain for DT is subject to both temporal and spatial background noise since the impact of drought depends importantly on micro-climates, agronomic features such as soil types, and the yield impact of the timing and persistence of rainfall shortages. Thus, the relative DT yield gain is $dz_{DT}\,(x_t) = z_{DT}(x_t) - z(x_t) + \varepsilon_z$ where $\varepsilon_z \sim N(0,\sigma_z)$. These relative yield gains translate into relative utility benefits:

$$du_{Bt}(x_t) = u[\bar{z}_{Bt} - c, \bar{w}(1-\delta_{Bt}x)] - u[z(x_t), \bar{w}(1-\delta_{Bt}x)]$$

Table 2. Interaction and adoption probabilities for baseline simulation.

Interaction	Pr(Interaction)	Pr(choose 0)	Pr(choose 1)
A_1, A_1	$(1-r)p^2$	0	1
A_1, A_0	$r + 2p(1-p)(1-r)$	$\frac{1}{2} + \beta(u_{0,t-1} - u_{1,t-1})$	$\frac{1}{2} + \beta(u_{1,t-1} - u_{0,t-1})$
A_0, A_0	$(1-r)(1-p)^2$	1	0

$$du_{DT}(x_t) = u[z_{DT}(x_t), \overline{w}(1-\delta_{DT}x)] - u[z(x_t), \overline{w}(1-\delta_{DT}x)] + \varepsilon_u, \quad (2)$$

where $\varepsilon_u \sim N(0,\sigma_u)$ reflects the random noise ε_z. Farmers in this model learn about the merits of the new trait on the basis of these utility benefits ($du_i(x_i)$) computed from the relative yield gains ($dz_i(x_i)$) they observe in their own plots or others' fields through social interaction. Figure 1 depicts relative yield dz and relative utility du as a function of x.[2] The relative utility benefit initially increases more rapidly for DT than Bt because $\delta_{Bt} < \delta_{DT}$, which makes the relative yield gain from DT especially valuable. The general sensitivity of income to drought is not universally good for the relative utility benefit of DT however: during extreme drought—precisely when individuals most need some relief from drought stress—DT benefits can disappear. The right panels of Figure 1 depict our construction of the probability distribution of x (top right) using a weighted composition of two beta distributions (bottom right), which allows us to control the frequency of extreme events by adjusting h.

Interaction and Diffusion

To model diffusion of the new traits, we allow individuals to interact and learn from each other (see Bowles [2004] and Henrich [2004] for similar interaction and learning models). In this formulation, farmers' individual adoption decisions are simplistic, but the diffusion process is dynamic and follows evolutionary principles. We therefore focus more on diffusion dynamics at the population level than on atomistic adoption decisions. At this population level, the evolutionary concept of fitness functions as the analogue to an individual level utility function (see Boyd & Richerson, 1985; Richerson & Boyd, 1993).

Farmers in our model observe the crop yields of those with whom they interact. They then use this relative yield gain and their own utility function to determine the relative utility benefit of the new trait, which

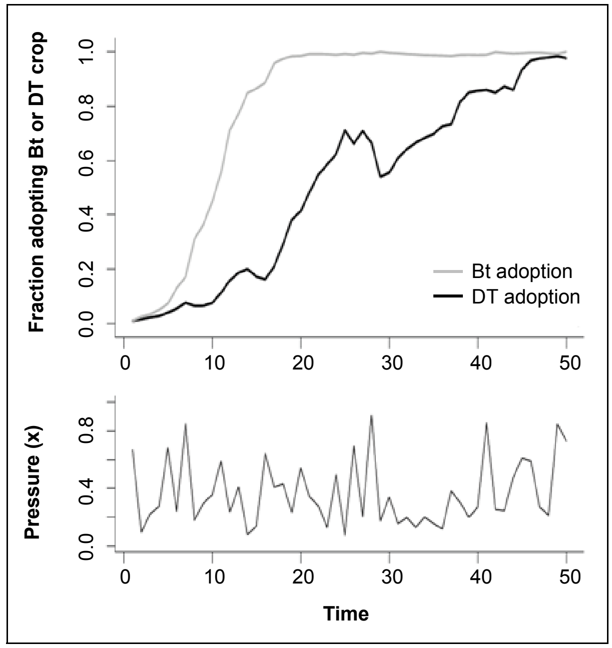

Figure 2. A single simulation of Bt versus DT diffusion assuming parameter values from baseline scenario: $\overline{z}$=1, $\overline{z}_{Bt}$=1.2, $\overline{z}_{DT}$=1, $\overline{w}$ = 4, δ_{Bt} = 0.1, δ_{DT} = 0.6, α = 2, x_1 = 0.5, x_2 = 0.8, c = 0.1, σ_u = 0.1, β = 10, h = 0.2.

remains the basis for learning about and adopting the new trait. Specifically, we let A_0 and A_1 represent adopters of the baseline variety 0 and variety 1 with the new trait, respectively, and p be the fraction of the population that are A_1. The interaction and adoption probabilities are shown in Table 2, where β is a constant that determines how utility differentials affect the probability of adoption and r is a probability adjustment that indicates the likelihood that A_0 and A_1 interact (when r=0, this interaction is random). These interaction and adoption probabilities allow us to express the dynamics of diffusion as a recursion function on p:

$$P_{t+1} = [r + 2p_t(1-p_t)(1-r)]\,[\tfrac{1}{2} + \beta(u_{1,t} - u_{0,t})] + (1-r)\,p_t^2 \quad (3)$$

This recursion function sets the stage for our baseline simulations that compare the diffusion paths of Bt and DT. To enable more detailed simulations of DT diffu-

2. *The baseline parameter values used for this figure are as follows:* $\overline{z}$=1, $\overline{z}_{Bt}$=1.2, $\overline{z}_{DT}$=1, $\overline{w}$=4, δ_{Bt}=0.1, δ_{DT}=0.6, α=2, x_1=0.5, x_2=0.8, c=0.1, σ_u=0.1

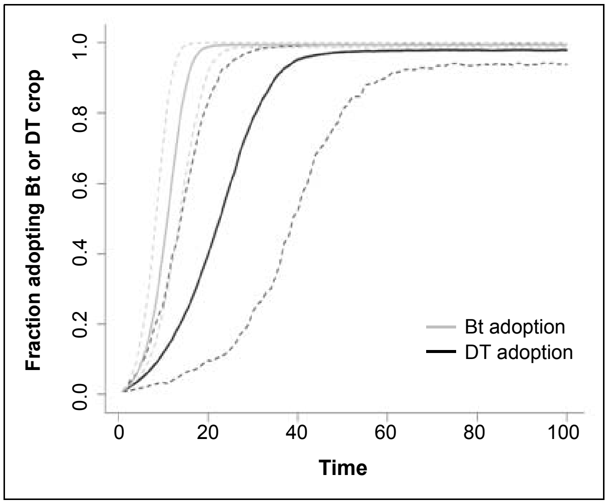

Figure 3. Median (solid) and 95% confidence intervals (dashed) for Bt and DT diffusion for 1,000 runs.

sion, we modify this recursion function below to allow for heterogeneous farmers.

DT research is frequently motivated by the goal of reducing drought vulnerability of the poor. Implicit in this motivation is the idea that DT varieties will be particularly valuable to vulnerable farmers who are particularly risk averse. To illustrate how risk aversion shapes DT diffusion in our model, we introduce two types of individuals: high risk-averse types with α_H=2 (as above) and low risk-averse types with α_L=0.2. This requires us to specify a new set of probabilities of interaction and adoption in order to construct new recursion function. The Appendix shows the derivation of recursion functions for heterogeneous farmers.

Diffusion Simulations

Baseline Bt versus DT Simulation

We simulate several consecutive seasons by drawing x from the composite probability distribution in Figure 1 and using the yield, utility, and recursion functions above to track diffusion of the new trait. We simulate Bt and DT separately and then compare diffusion paths for the two in Figure 2. To be clear, this simulation does not model the decision between Bt and DT, but rather between the baseline variety and Bt variety and—in a separate simulation—between the baseline variety and DT variety. Even with a positive seed premium on Bt, diffusion is nearly complete 15 years after release. Diffusion of DT is much more sluggish (even without a seed premium) primarily because its relative benefits are not monotonically increasing in drought pressure

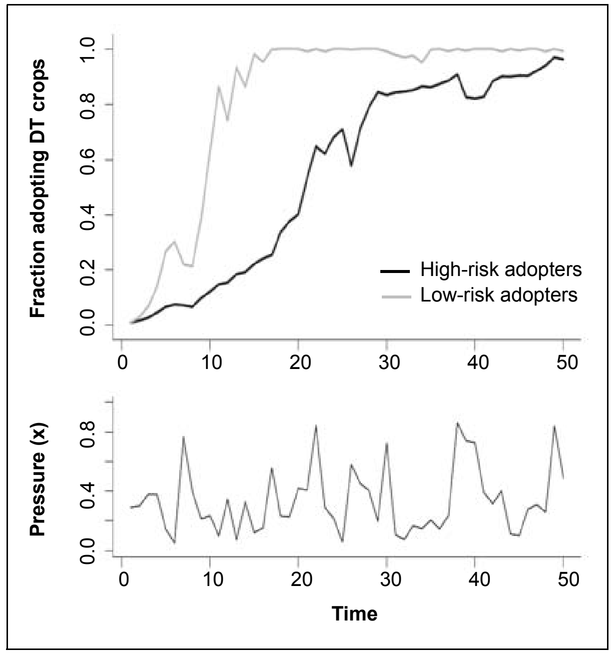

Figure 4. A single simulation of DT diffusion for high and low risk averse farmers for parameter values as in Figure 2 with α_H = 2, α_L = 0.2, *r* = 0.1.

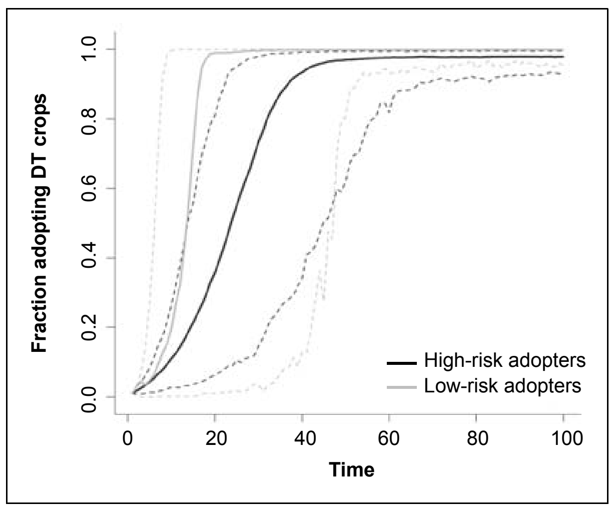

Figure 5. Median (solid) and 95% confidence intervals (dashed) for DT diffusion with high risk averse and low risk averse farmers for 1,000 runs.

(e.g., disadoption in season 29 is due to extreme drought pressure of x=0.85 in the prior season). To capture expected diffusion paths, we compute the median and 95% confidence interval for Bt and DT based on 1,000 simulations (Figure 3). Diffusion of DT in this stylized model is expected to take four times longer than Bt.

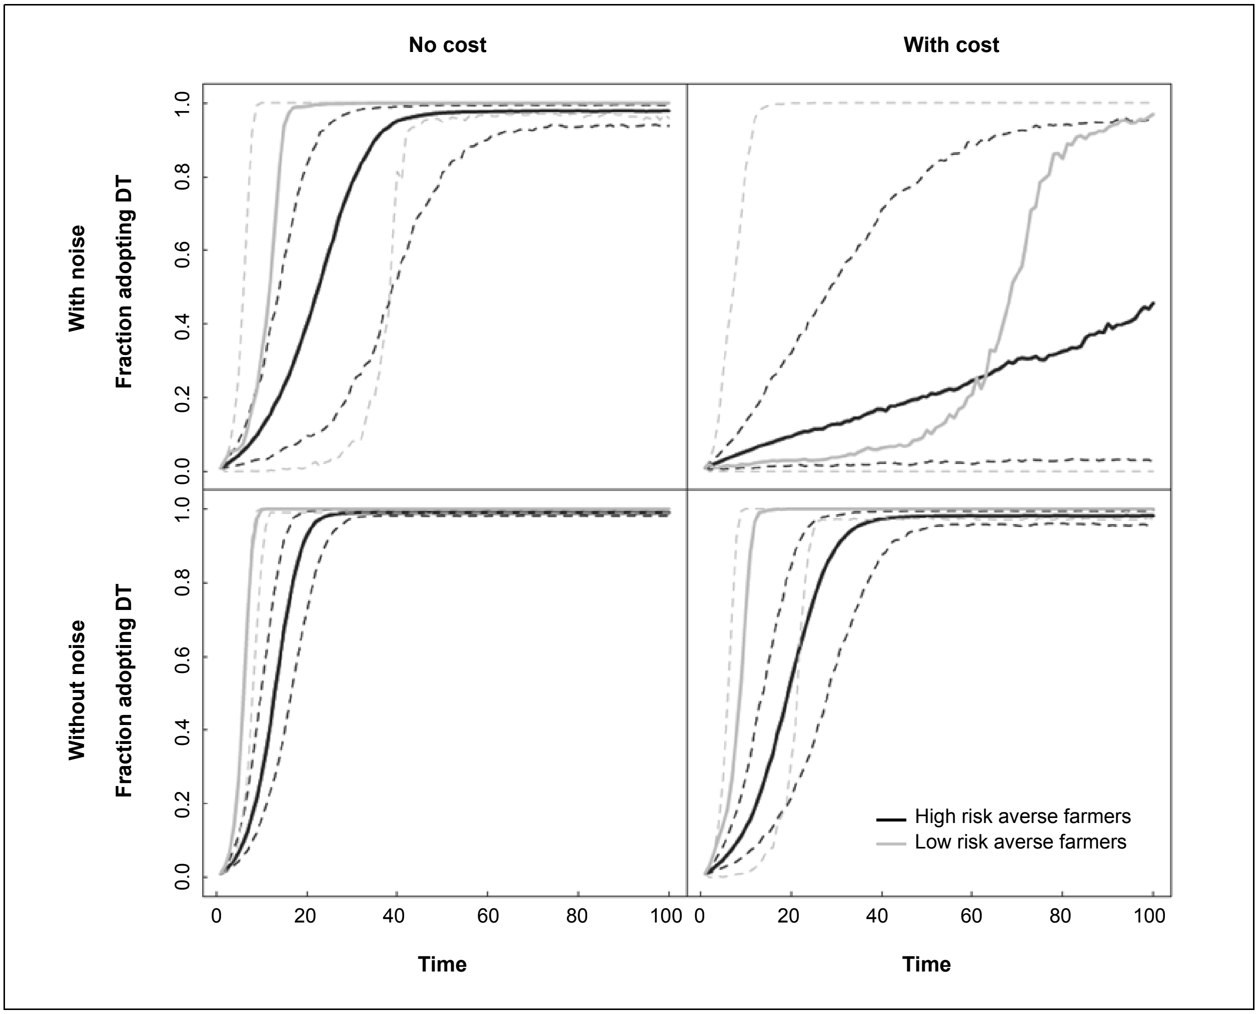

Figure 6. High and low risk averse farmers with and without background noise and additional seed cost for DT.

Heterogeneous Risk Aversion and DT Diffusion Simulation

To demonstrate the impact of risk aversion (which is often cited as a motivation for reducing drought vulnerability of the poor) on DT diffusion, we compare adoption and diffusion of high and low risk-averse farmer types using the recursion functions derived in the Appendix. The simulated DT diffusions with these risk-averse types are shown in Figures 4 and 5 (henceforth we focus on DT diffusion exclusively). While some sequences of drought pressure can indeed lead high risk-averse farmers to adopt more rapidly than low risk-averse types, diffusion among the high risk-averse types tends to take three times longer. These farmers could no doubt benefit from DT protection when the right kind of drought occurs, but DT can also offer no protection at all under extreme drought. Highly risk-averse farmers are especially sensitive to this lack of protection when they most need it. The variable relative return to DT slows the adoption of DT among high risk-averse farmers relative to their low risk-averse counterparts.

Background Noise and Seed Cost Simulation

Next, we inspect the effect of background noise and seed cost on DT diffusion by farmer type. The volume of background noise in the learning process is directly correlated with the marginality of a farmer's production environment. Since marginal farmers are an oft-cited clientele of DT research, understanding how background noise affects DT adoption and diffusion is important. Similarly, concerns about getting new DT seeds in the hands of the poor, along with initially high seed premia among Bt crops, have put seed pricing on the table at early stages of DT research, particularly in public-private DT partnerships. Additional diffusion

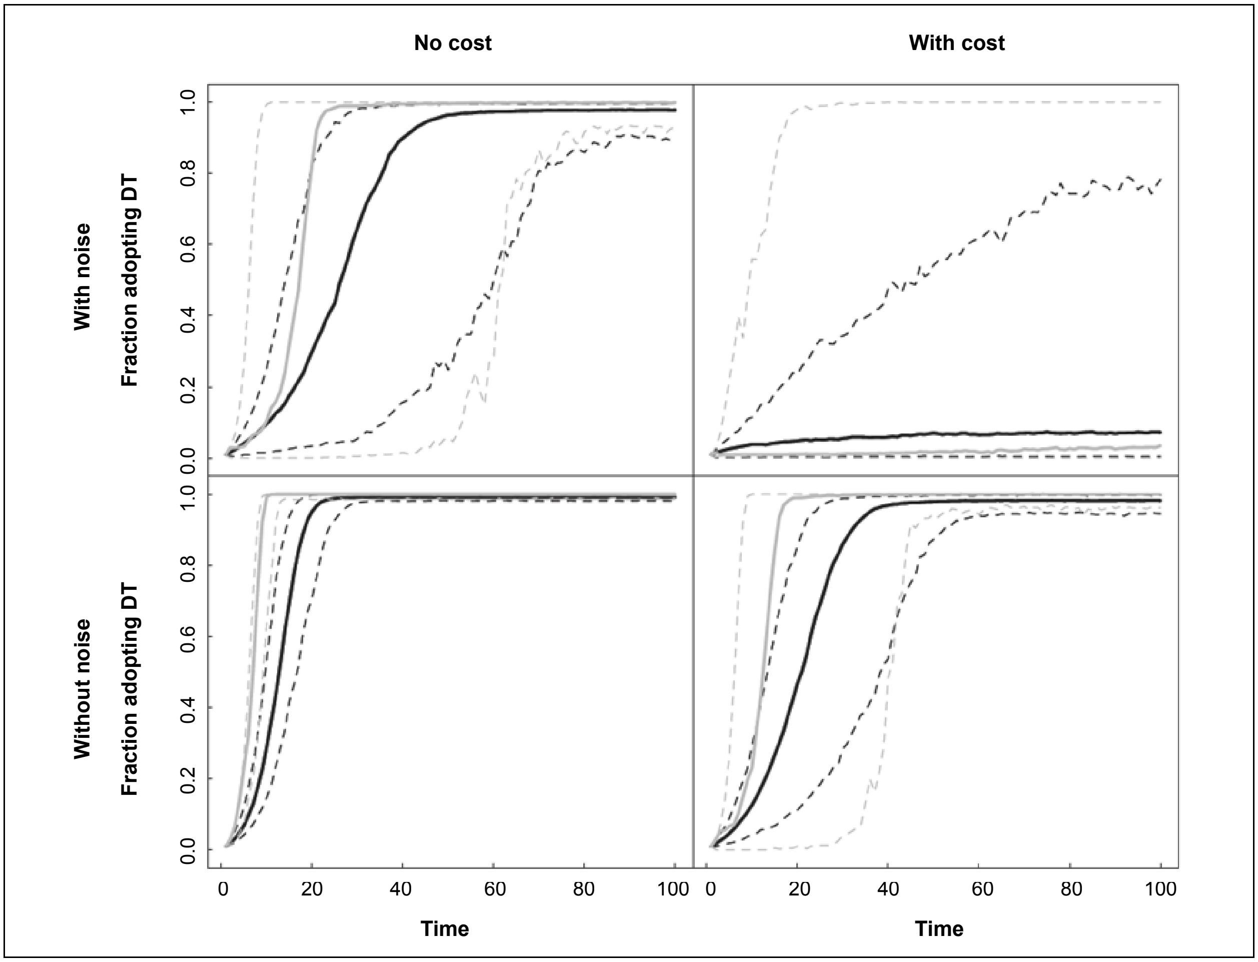

Figure 7. Climate change scenario (*h*=0.4) for high and low risk averse farmers.

scenarios without background noise and with a positive seed cost premium are shown in Figure 6. Charging a positive seed cost (c=0.1) for DT in the presence of background noise (upper-right panel) devastates diffusion, although low risk-averse farmers ultimately learn the potential value of DT and adopt quickly—on average around season 60. Interestingly, the comparison of the effect of noise and a positive seed premium in isolation (upper-left and lower-right, respectively) suggests that the marginality of growing conditions and the associated background noise would slow learning and diffusion more than charging a premium for DT would. As before, where there appears to be a difference, high risk-averse farmers tend to adopt more slowly than their low risk-averse peers.

Climate Change Simulation

The final commonly cited motivation for DT research we address is impending climate change, which in many parts of the world is expected to reduce rainfall and increase the probability of extreme weather events, including drought. We model the impact of increasing the probability of extreme drought events in our simple simulation by increasing h to 0.4 (recall Figure 1). Capturing the complexity of climate change in this stylized way is obviously abstract, but it is nonetheless illustrative. We run this simulation with and without background noise and additional seed cost as above and depict the results in Figure 7. The overall effect of this stylized climate change is to make DT diffusion more erratic and less predictable (i.e., to widen the confidence intervals). This occurs because a greater frequency of extreme drought events makes it more difficult for farmers to observe the merits of the DT. Indeed, with a positive seed cost farmers experience a relative loss during these events, hence the collapse of diffusion paths in the upper-right panel.

Two notes about climate change effects are worth making. First, since these effects are based on the comparison of Figures 5 and 6, they are conditioned on our baseline assumptions (e.g., h=0.2). In locations with infrequent, mild droughts initially (e.g., h=0.0), climate change may actually speed learning and adoption. Thus, we do not claim that climate change will universally slow learning and diffusion of DT varieties, but rather that it will not facilitate the process either in many arid and semi-arid locations. Second, consider a comparison between climate change effects on the diffusion of DT varieties and on the uptake drought insurance. Since drought insurance confers monotonically increasing (or at least non-decreasing) benefits as drought severity as measured by a rainfall index increases, climate change would speed rather than slow the diffusion of such an insurance product because extreme drought events exaggerate the value of being insured relative to not being insured and therefore maximize the information content of a given season.

Discussion and Conclusion

The potential welfare benefits of DT are clear and compelling in the context of rainfed agriculture in arid and semi-arid parts of developing countries. We do not dispute these potential benefits but caution against a 'build it and they will come' approach that ignores the practical complications farmers will face in assessing the relative merits of DT varieties. We argue that adoption among vulnerable, risk-averse farmers who face marginal growing conditions may be particularly sluggish. More frequent extreme drought events complicate rather than facilitate learning and DT adoption among this common target clientele. Concerns about DT seed pricing appear to be well-founded: even a modest additional seed cost dramatically slows diffusion.

Where might the assumptions of our stylized model mislead or misrepresent? The core assumption of the model—the non-monotonicity of relative DT benefits—seems robust even if the specific parameterization is debatable. Still, a more realistic relative benefits profile is unlikely to change the results much as long as the possibility of no meaningful relative benefits under extreme drought still exists. Our probability distribution of drought severity is a simplification in many ways. Drought is a complex event that has crucial timing and agronomic (e.g., soil types and slopes) dimensions. Further, our representation of climate change is admittedly crude and intended only to illustrate the complications that more frequent extreme drought could introduce.

Among our assumptions, two are especially worthy of additional discussion. First, we assume that DT does not confer an unconditional relative yield benefit. While this seems defensible in cases where the focus is on limiting drought losses, in settings where water availability nearly always constrains crop productivity, generalized gains in water-use efficiency may confer relative benefits across a broader range of rainfall outcomes. This could facilitate learning and speed diffusion. Although DT research will never fully overcome the non-monotonicity of relative benefits problem since at least some water is required to grow anything, achieving generalized improvements in water-use efficiency will remain an important research objective in arid and semi-arid regions with ever-binding water availability constraints. In a similar vein, early maturation—which can convey drought benefits by reducing exposure to late-season precipitation—may facilitate learning: although the relative benefits from such a trait are conditional on late season rainfall, its expression is at least observable every season. Second, our model does not address spatial dimensions to DT adoption and diffusion. This is a necessary simplification, but it glosses over what might be an important dynamic response: the effect DT will have on broadening the zones in which a particular crop is grown. Indeed, this change in the allocation crop production across space may be one of the most important impacts of DT crops.

While DT crops do confer quasi-insurance benefits, the terms of any implicit insurance policy contained in DT crops are concealed from farmers, who must discover them from experience. In conjunction with the non-monotonicity of relative DT benefits, these concealed terms of drought protection have important implications for behavioral household responses. Whereas a household with an explicit drought insurance policy may indeed become more entrepreneurial in other parts of its livelihood strategy (so-called "positive moral hazard"), a household adopting DT crops is much less likely to respond in the same way because the terms of protection are neither explicit, certain, nor do they provide much protection to extreme drought events. DT may look like drought insurance in the right kind of drought, but households are unlikely to reformulate their livelihood portfolios in response to such fickle insurance protection.

References

Bowles, S. (2004). *Microeconomics: Behavior, institutions, and evolution.* Princeton, NJ: Princeton University Press.

Boyd, R., & Richerson, P.J. (1985). *Culture and the evolutionary process.* Chicago, IL: University of Chicago Press.

Conley, T.G., & Udry, C.R. (2010). Learning about a new technology: Pineapple in Ghana. *American Economic Review*, *100*(1), 35-69.

Foster, A-D., & Rosenzweig, M-R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, *103*(6), 1176-1209.

Henrich, J. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior and Organization*, *53*(1), 3-35.

Herring, R.J. (2008). Whose numbers count? Probing discrepant evidence on transgenic cotton in the Warangal district of India. *International Journal of Multiple Research Approaches*, *2*(2), 145-159.

Marra, M., Pannell, D.J., & Abadi, G.A.K. (2003). The economics of risk, uncertainty and learning in the adoption of new agricultural technologies: Where are we on the learning curve? *Agricultural Systems*, *75*(2/3), 215-234.

Pray, C., Ma, D., Huang, J., & Qiao, F. (2001). Impacts of Bt cotton in China. *World Development*, *29*(5), 813-825.

Qaim, M. (2003). Bt cotton in India: Field trial results and economic projections. *World Development*, *31*(12), 2115-2127.

Qaim, M., Subramanian, A., Naik, G., & Zilberman, D. (2006). Adoption of Bt cotton and impact variability: Insights from India. *Review of Agricultural Economics*, *28*(1), 48-58.

Qaim, M., & Zilberman, D. (2003). Yield effects of genetically modified crops in developing countries. *Science*, *299*(5608), 900-902.

Pennisi, E. (2008, April 11). Plant genetics: The blue revolution, drop by drop, gene by gene. *Science*, *320*(5873), 171-173.

Richerson, P., & Boyd, R. (1993). Rationality, imitation, and tradition. In R.H. Day & P. Chen (Eds.), *Nonlinear dynamics and evolutionary economics.* New York: Oxford University Press.

Smale, M., Zambrano, P., Gruère, G., Falck-Zepeda, J., Matuschke, I., Horna, D., et al. (2009). *Measuring the economic impacts of transgenic crops in developing agriculture during the first decade: Approaches, findings, and future directions* (Food Policy Review 10). Washington, DC: International Food Policy Research Institute (IFPRI).

Subramanian, A., & Qaim, M. (2009). Village-wide effects of agricultural biotechnology: The case of Bt cotton in India. *World Development*, *37*(1), 256-267.

Thirtle, C., Beyers, L., Ismael, Y., & Piesse, J. (2003). Can GM-technologies help the poor? The impact of Bt cotton in Makhathini Flats, KwaZulu-Natal. *World Development*, *31*(4), 717-732.

Tripp, R.B. (2009). *Biotechnology and agricultural development: Transgenic cotton, rural institutions and resource-poor farmers.* New York: Routledge.

Acknowledgements

We thank Steve Vosti, Carl Pray, Guillaume Gruère, and seminar audiences at UC Davis, the January 2009 "Biotechnology in Developing Countries" conference at UC Berkeley, and the 2010 Pacific Development Conference at USC for helpful comments and suggestions.

Appendix

Derivation of Recursion Functions for Heterogeneous Farmers

Adding high and low risk-averse types expands the recursion function that determines how farmers interact and how they learn from each other. A fraction p_H are high-risk adopters choosing technology 1, and are represented by H_1, and the other fraction $(1 - p_H)$ are high-risk adopters (H_0) choosing technology 0. Similarly, a fraction p_L and $(1 - p_L)$ are low-risk adopters choosing technology 1 (L_1) and 0 (L_0), respectively. The utility of either technology 0 or 1 to the high-risk adopter becomes

$$u_0^{(H)} = \frac{[w(1 - \delta x_t) + z(1 - x)]^{1 - \alpha_H}}{1 - \alpha_H}$$

$$u_1^{(H)} = \frac{[w(1 - \delta x_t) + \bar{z} - c]^{1 - \alpha_H}}{1 - \alpha_H}$$

and for low-risk adopters is

$$u_0^{(L)} = \frac{[w(1 - \delta x_t) + z(1 - x)]^{1 - \alpha_L}}{1 - \alpha_L}$$

$$u_1^{(L)} = \frac{[w(1 - \delta x_t) + \bar{z} - c]^{1 - \alpha_L}}{1 - \alpha_L} \quad .$$

Both high- and low-risk adopters learn from each other and themselves about the utility of either technology 1 or 0, influencing their choice of technology in the next season. A fraction k of the time, adopters' interaction with individuals that share the same risk aversion, it follows then that $(1 - k)$ of the time individuals interact with another with a different level of risk aversion. As above, r is the fraction of the time from random that individuals view the performance of the other technology. Table A1 details the interactions among adopters

Table A1. Interaction and adoption probabilities with high and low risk-averse farmer types.

Self	Other	Pr(Other|Self) × Pr(Self)	Pr(H_1|Other)	Pr(H_0|Other)	Pr(L_1|Other)	Pr(L_0|Other)
H_1	H_1	$k(1-r)p^2_H$	1	0	0	0
H_1	H_0	$k[r+(1-r)(1-p_H)]p_H$	$½+\beta(u_1^{(H)}-u_0^{(H)})$	$½+\beta(u_0^{(H)}-u_1^{(H)})$	0	0
H_0	H_1	$k[r+(1-r)p_H](1-p_H)$	$½+\beta(u_1^{(H)}-u_0^{(H)})$	$½+\beta(u_0^{(H)}-u_1^{(H)})$	0	0
H_0	H_0	$k(1-r)(1-p_H)(1-p_H)$	0	1	0	0
L_1	L_1	$k(1-r)p^2_L$	0	0	1	0
L_1	L_0	$k[r+(1-r)(1-p_L)]p_L$	0	0	$½+\beta(u_1^{(L)}-u_0^{(L)})$	$½+\beta(u_0^{(L)}-u_1^{(L)})$
L_0	L_1	$k[r+(1-r)p_L](1-p_L)$	0	0	$½+\beta(u_1^{(L)}-u_0^{(L)})$	$½+\beta(u_0^{(L)}-u_1^{(L)})$
L_0	L_0	$k(1-r)(1-p_L)(1-p_L)$	0	0	0	1
H_1	L_1	$(1-k)(1-r)p_Lp_H$	1	0	0	0
H_1	L_0	$(1-k)[r+(1-r)(1-p_L)]p_H$	$½+\beta(u_1^{(H)}-u_0^{(H)})$	$½+\beta(u_0^{(H)}-u_1^{(H)})$	0	0
H_0	L_1	$(1-k)[r+(1-r)p_L](1-p_H)$	$½+\beta(u_1^{(H)}-u_0^{(H)})$	$½+\beta(u_0^{(H)}-u_1^{(H)})$		
H_0	L_0	$(1-k)(1-r)(1-p_L)(1-p)$	0	1	0	0
L_1	H_1	$(1-k)(1-r)p_Hp_L$	0	0	1	0
L_1	H_0	$(1-k)[r+(1-r)(1-p_H)]p_L$	0	0	$½+\beta(u_1^{(L)}-u_0^{(L)})$	$½+\beta(u_0^{(L)}-u_1^{(L)})$
L_0	H_1	$(1-k)[r+(1-r)p_H](1-p_L)$	0	0	$½+\beta(u_1^{(L)}-u_0^{(L)})$	$½+\beta(u_0^{(L)}-u_1^{(L)})$
L_0	H_0	$(1-k)(1-r)(1-p_H)(1-p_L)$	0	0	0	1

and the probability of choosing a technology. From this table, we can derive the adoption recursion for high-risk adopters as

$$
\begin{aligned}
P_{H,t+1} = {} & (1-r)\,p_{H,t}\,[(1-k)\,p_{L,t} + kp_{H,t}] \\
& + [½ + \beta(u_1^{(H)} - u_0^{(H)})] \\
& \times (r + (1-r)\,\{[(1-k)\,(1-r)\,p_{L,t} \\
& + p_{H,t}\,[1 + \mathrm{k}(1-2p_{H,t}) - (1-k)2p_{L,t}]\})
\end{aligned}
$$

and low-risk adopters as

$$
\begin{aligned}
p_{L,t+1} = {} & (1-r)\,p_{L,t}\,[(1-k)\,p_{H,t} + kp_{L,t}] \\
& + [½ + \beta(u_1^{(L)} - u_0^{(L)})] \\
& \times (r + (1-r)\,\{[(1-k)p_{H,t}(1-2p_{L,t}) \\
& + p_{L,t}\,[1 + \mathrm{k}(1-2p_{L,t})]\}).
\end{aligned}
$$