Flock Management for Solar Grazing

Vincent W. Slaugh
SC Johnson College of Business, Cornell University, Ithaca, New York 14853, USA, vslaugh@cornell.edu

Linwei Xin
School of Operations Research and Information Engineering, Cornell University, Ithaca, New York 14853, USA, lx267@cornell.edu

Shane G. Henderson
School of Operations Research and Information Engineering, Cornell University, Ithaca, New York 14853, USA, sgh9@cornell.edu

March 9, 2026

To control vegetation, solar asset owners employ farmers to graze sheep on solar sites. While earning revenue from vegetative maintenance, farmers also have the opportunity to provide agricultural value through raising and selling sheep. By accounting for flock dynamics of births, deaths, and sales, we model the operational economics of solar grazing for farmers and focus on key flock management decisions.

We model solar grazing operations as a finite-horizon linear program that jointly optimizes the initial procurement decisions and subsequent sales decisions, subject to age-structured flock dynamics and a target flock-size constraint. During the flock growth phase, the “breed-versus-buy” decision involves weighing the upfront costs and anticipated revenue from lamb and sheep sales against heightened costs of vegetation management when the flock size is below its target level. By analyzing this trade-off, we demonstrate that the farmer’s total profit function is piecewise-linear and concave in the number of sheep initially purchased externally. Beyond the flock growth phase, we prove that an optimal steady-state solution exhibits a maximum-lifetime structure: ewes are retained up to a maximal age, and sales are restricted to lambs and final-year ewes. Finally, we extend the model to incorporate stochastic births and deaths, showing that a linear programming approach yields a tight bound on achievable performance in the stochastic model for large flocks and long time horizons, and provide a practical management policy that performs strongly in numerical examples.

Our results help small-scale farmers, renewable energy developers, and policymakers understand how solar grazing enhances the environmental benefits of solar farms while preserving agricultural value on land converted to solar use. In partnership with a solar grazing platform that supports solar grazers and connects them with site owners, a spreadsheet implementation of our model has helped dozens of farmers make procurement decisions, predict cash flows, and even secure financing to purchase a flock.

Key words: sustainability; service operations; agricultural economics; agrivoltaics; perishable inventory

---

1. Introduction

The global surge in energy produced via solar farms faces a mundane but nevertheless critical operational challenge: mowing the grass. In addition to violating key safety and aesthetic standards for solar sites, overgrown vegetation can also reduce the efficiency of photovoltaic panels by blocking sunlight. One possible solution matches the modern technology of solar panels with the ancient practice of tending sheep. By grazing on vegetation, sheep eliminate the need for labor-intensive landscaping and gas-powered equipment, whose emissions undermine the environmental benefits of solar energy generation. Using this vegetation as a food source, lambs and sheep can grow and provide meat that would otherwise require dedicated pastureland to produce.

The rise of solar energy has been remarkable, with global photovoltaic installations adding over 400 gigawatts in 2023 and bringing the total cumulative capacity to 1.6 terawatts. In the United States, solar energy contributed 54% of all new electric generation capacity in 2023, up from just 6% in 2010 (Feldman et al. 2024). By 2040, solar farms could occupy up to 7 million acres in the U.S., with about half of that land expected to be prime farmland (Sorensen et al. 2022). This shift in land use raises important public policy questions and has sparked strong opposition to solar development in many rural communities, partly due to concerns about the loss of food production capacity. Coupling solar site development with sheep grazing offers a promising solution.

Operating a solar site requires effective vegetation management to prevent the shading of photovoltaic arrays and interference with equipment. Because goats may damage equipment and cows require costly modifications to raise the arrays, sheep are the most viable option for widespread solar grazing. Historically, sheep were raised primarily for wool, with meat as a byproduct. However, global trends have shifted the focus to meat. Given the small size of the wool industry in the United States, "hair breeds" of sheep that do not require shearing and are raised for meat predominate in solar grazing operations across much of the country. Flock management decisions are particularly critical in the U.S., where the sheep population has declined from 56 million in 1942 to around 5 million today (Thorne et al. 2021). Mazinani and Rude (2020) estimate a more robust global sheep population of 1.173 billion. Solar grazing is a global phenomenon; for example, Bowen (2024) describes its implementation in Australia and predicts that meat sheep will become the predominant grazing animal due to the increased costs of setting up solar sites to accommodate beef steers.

For common meat breeds of sheep in the northeastern United States, ewes typically have a productive lifespan of 5-6 years and typically give birth to twins or single lambs each winter or spring. Most meat sold as "lamb" in the United States comes from sheep under 14 months old, conveniently aligning with religious holidays when lamb is traditionally consumed. Meat from older sheep has much lower domestic demand but can be exported or supplied in the pet food industry

(USDA Economic Research Service 2025). Sheep are often sold in lots at auctions, and the U.S. Department of Agriculture Agricultural Marketing Service (2025) provides weekly price reports.

Beyond the scope of this paper, solar grazing and other agricultural activities on solar farms — collectively known as "agrivoltaics" — offer significant co-benefits for livestock, farmers, and rural communities. Fonsêca et al. (2023) demonstrate that shade from solar panels enhances both animal welfare and productivity, reducing wool surface temperatures by 7-8°C. With stable income streams from maintenance contracts and relatively low capital requirements, solar grazing also provides rural communities with a viable pathway to more attractive and sustainable agricultural livelihoods. Furthermore, enclosed solar sites offer opportunities to promote biodiversity through initiatives such as pollinator-friendly vegetation and bird boxes. Notably, our research partner contributed to efforts to reintroduce bobwhite quail — extirpated from much of the northeastern United States due to habitat loss — by using solar sites as restored habitat.

Farmers who enter into solar grazing contracts must make operational decisions about their sheep flocks that blend elements of workforce and inventory management. Effective decision-making involves balancing the initial capital costs of acquiring a flock with ongoing revenues from grazing and lamb sales, while accounting for potential losses from predation or suboptimal culling, as well as labor costs for vegetation management if the flock is too small to fully graze the site.

As the first to study solar grazing from an operations management perspective, we develop novel operational models to investigate optimal flock management decisions that balance vegetation control requirements, livestock dynamics, and profitability. We divide flock management into two stages: initial growth and long-term maintenance, and provide the following structural results:

1. At the start of operations, a key question is whether a farmer should purchase an initial flock or allow the flock to grow gradually through breeding up to a target flock size. This "breed-versus-buy" decision involves weighing the upfront costs and anticipated revenue from lamb and sheep sales against the ongoing costs of vegetation management when the flock size is below its target level. We analyze this tradeoff and show that the farmer's total profit function is piecewise linear and concave in the number of sheep initially purchased externally. For realistic parameter values, we find that it is typically optimal to purchase the entire target flock up front, rather than build it gradually through breeding. This result is quite robust, though certain input conditions can make gradual breeding more advantageous.
2. For long-term flock management, we show that incorporating complex multi-year cycles with varying age profiles provides no benefit. Moreover, under realistic assumptions, we prove that

the optimal steady-state solution exhibits a maximum-lifetime structure: ewes are retained up to a maximal age, and sales are restricted to lambs and final-year ewes. This maximum-lifetime policy aligns with a widely accepted heuristic; see, e.g., OSU Sheep Team (2018).

3. Extending our flock maintenance model to incorporate stochastic births and deaths, we establish the intuitive result that no policy can asymptotically outperform the long-run profit of the corresponding deterministic model. We further propose a policy that provably asymptotically achieves the deterministic model's optimal performance as the flock size grows. In other words, random fluctuations do not erode the core insights from the deterministic model. These results are reminiscent of a standard result in revenue management that a certain linear program delivers a tight upper bound on the performance of revenue-management policies (Talluri and Van Ryzin 2006, Chapter 3), although the proof techniques are fundamentally different. We reinforce the practical value of the deterministic model by adapting it to support sales decisions within a model-predictive control policy that exhibits impressive empirical performance.

In addition to our theoretical results, a spreadsheet implementation of the breed-versus-buy decision has been incorporated into the training materials of a solar grazing network, supporting dozens of farms and helping some farmers secure loans to procure flocks. Consultants have also leveraged the spreadsheet to assess contracts and provide operational guidance for utility-scale operations. Furthermore, the model — which quantifies the agricultural value of solar grazing — has been used to get solar sites approved in multiple municipalities in the northeast United States where town leaders had previously rejected authorizing solar sites due to a fear of losing farmland.

The remainder of this paper is organized as follows. In Section 2, we review the related literature on solar grazing and connect it to themes in the operations management literature. Section 3 models the flock growth phase over a finite horizon as a linear program that jointly optimizes the initial procurement decisions and subsequent sales decisions. The resulting insights inform the implementation of our model by a platform that supports farm partners in becoming solar grazers and connects them with solar site owners. In Section 4, we theoretically analyze the performance of the maximum-lifetime policy in steady-state. Section 5 investigates a stochastic model of flock dynamics, incorporating stochastic births and deaths. There we show that 1) no operating policy can substantially outperform the deterministic solution, at least as flock size increases; 2) the performance of a tailored policy is provably near-optimal for large flock sizes; and 3) by way of numerical results, a simple operating policy based on the solution to a linear programming formulation appears to be very close to optimal over a wide range of flock sizes. We conclude and discuss potential future research questions in Section 6. All proofs are deferred to the e-companion.

2. Literature Review

By analyzing a maintenance solution for solar sites, our work responds to the call by Secomandi (2025) to study the clean energy transformation at the asset level and provide actionable insights for asset owners, operators, and investors. In modeling sheep grazing as a viable solution for solar sites, we build on the operational challenges faced by "cleantech" firms as documented by Plambeck (2013). To date, utility-scale solar site operations have received very limited attention. Wang et al. (2024) propose an integer program to optimize the layout of cables and other key components of a solar site. More broadly, Peng et al. (2024) analyze how renewable energy sources, such as solar, integrate with existing energy production and storage systems from a policy perspective.

Methodologically, the problem of flock management is closely related to classical stochastic perishable inventory management (e.g., Karaesmen et al. (2011), Nahmias (2011), Chen et al. (2014), Chao et al. (2015), Bu et al. (2023)). In such traditional settings, the inventory state is represented by the number of items at different ages (for example, blood of various ages), and decision-makers determine replenishment and disposals after demand is realized in each time period. While our flock management problem shares some similarities with this framework (such as aging items with finite lifespan), it also exhibits distinct characteristics. Specifically, in our setting, inbound inventory is primarily determined by births within the existing flock, and the decision-making centers on when to sell. Moreover, unlike traditional inventory models, there is no demand to satisfy; instead, the objective is to maintain the flock size at a desired level.

Solar grazing has already attracted attention from agricultural scientists and economists. Kochendoerfer et al. (2019) report data from grazing sheep on a 22-acre solar site in upstate New York, estimating a reduction in annual landscaping labor from 344 hours to 139 hours. They also survey operations and maintenance managers in 2018 and estimate a cost of $828 per acre for vegetation management on solar sites that did not utilize grazing. Kochendoerfer and Thonney (2021) follow up with another study assessing the potential for solar grazing to scale in New York State. Shiflett (2022) provides detailed budget projections and cost analysis for a 1,060-acre site in Mount Morris, NY, based on buying 9,540 lambs each spring to be sold in the fall. Gasch et al. (2025) analyze both small- and large-scale solar grazing operations with sheep and conclude that vegetation maintenance contracts make solar grazing more profitable than typical agricultural operations, especially when sheep are procured at auction. Beckwith (2021) provides a spreadsheet tool to support grazing decisions when solar sites are divided into multiple paddocks.

Our work's focus on the flock growth trajectory addresses a management challenge that has received relatively little attention from agricultural economists. In early work modeling livestock in Africa, Dahl and Hjort (1976) provide numerical analysis of how different parameters, such as life expectancy, affect livestock population trajectories. However, most existing literature focuses on management decisions after the herd or flock has reached a steady-state population, which aligns more closely with the second of our contributions on steady-state flock management. For example, in an effort to promote domestic sheep farming prior to the rise of solar grazing, Ontario Sheep Marketing Association (2012) provides an economic analysis of a flock assuming a steady-state population of ewes and lambs, without differentiating the age of ewes. Stygar et al. (2014) use a Markov decision process to study breeding decisions focused on calf sales, rather than herd growth or age-profile management across the entire herd. Similarly, Yan et al. (2022) develop a decision support tool for selling cattle before or after winter based on a price forecasting model, without considering age-specific herd policies. Farrell et al. (2020) simulate the effects of two breeding strategies and lambing rates for a representative New Zealand farm. In contrast, we analytically optimize over all possible breeding strategies to maintain a desired flock size.

Our research aligns with a broader stream of operations management work in agriculture that aims to support small farmers and to understand the incentives shaping their operational decisions. Ata et al. (2019) study how to effectively manage food and labor supply uncertainties in gleaning operations through a dynamic volunteer staffing policy. Yi et al. (2021) analyze how different financing strategies involving a capital-constrained smallholder farmer and an intermediary platform impact production and supply chain welfare. Agriculture is one application of the production and sales model focusing on keep-or-sell tradeoffs presented in Hu and Sun (2022) for "self-replicating" goods — a term which includes agricultural and animal husbandry goods. Chen and Ryan (2023) use a dynamic programming framework for the multi-year production planning problem of a farmer producing hops. Kouvelis et al. (2023) study a pork producer's livestock management decisions through the lens of commodity storage problems. Zhang et al. (2023) analyze contract farming using a game-theoretic model, highlighting the shortcomings of spot markets and showing that buyers with a long-term perspective are willing to pay higher prices to ensure farmers' viability. Levi et al. (2024) discuss their collaboration with the Indian government to implement an auction platform, demonstrating its positive impact on farmers' revenue. Velarde Morales and Xin (2026) use an optimization approach to design index-insurance contracts and assess the approach's impact on improving farmers' welfare using a dataset from a national insurer in Thailand.

3. Time-Dependent Flock Dynamics with Procurement

In this section, we develop a model of procurement and sales decisions that tracks the ewe population by age over time. We assume deterministic flock dynamics and relax the requirement that sheep counts be integers, allowing flock sizes to take nonnegative real values. This relaxation is reasonable when the flock is sufficiently large, as any rounding to whole animals has a negligible impact on aggregate outcomes. It also improves tractability and facilitates analytical insight.

Ewes are assumed to have a maximum integer lifespan of $n \geq 1$ years, at which point they are sold. Each age-i ewe produces exactly α_i ram lambs and α_i ewe lambs each year, for $i = 1, 2, \ldots, n$. Ram lambs are sold after one year. Ewe lambs may be sold before reaching one year of age; otherwise, they mature sexually in time to have their first lambs at one year old. A fraction ψ_i of age-i lambs and ewes are lost to disease and predation each year, for $i = 0, 1, 2, \ldots, n$. Sales decisions are made at the end of each year. Ewes cost c^H per year, which includes over-wintering costs and other miscellaneous costs such as veterinary work. Unless otherwise specified, we ignore upfront fixed costs and yearly fixed costs because they do not change the optimal decisions of how many ewes to procure and how to manage the flock. However, these costs *are* important when considering questions of overall profitability, so they are considered later in the section when we provide example results.

We assume the following sequence of events occurs each year:

1. New lambs are born at the start of the year.
2. Losses occur during the year, affecting both lambs and ewes.
3. Culling (equivalently, sales) is carried out before the end of the year.
4. The farmer incurs per head costs.

We now formulate an optimization problem that explores both initial procurement decisions and sales decisions over a finite time horizon of T years. Ram lambs are raised for one year before being sold at a price of p^R, which aligns with standard practice. Ewe lambs are either retained to replace lost or culled sheep, thereby maintaining the ewe population, or sold at a price p_0. At the end of their ith year, ewes can be sold at a price p_i, $i = 1, 2, \ldots, n$. A solar grazing operation also involves significant fixed costs (both at startup and annually), which we discuss in Section 3.2 but omit from this formulation.

The state and decision variables in this formulation are

1. $\mathbf{y}(t) \in [0, \infty)^n$ (recall that we relax integer restrictions): a vector representing the number of ewes of each age $1, \ldots, n$ at the *start* of year $t \in \{0, 1, \ldots, T\}$.

2. $\mathbf{z}(t) \in [0,\infty)^{n+1}$: a vector representing the number of ewes of each age to be sold during year $t \in \{0,1,\ldots,T-1\}$. Here, $\mathbf{z}_0(t)$ denotes the number of ewe lambs to be sold.
3. $\mathbf{x} \in [0,\infty)^n$ giving the number of ewes we initially procure, broken down by age $i = 1,2,\ldots,n$ years.

Recall that α_i denotes the reproduction rate and ψ_i denotes the loss rate for age-i ewes. The state dynamics are given, for each $t = 1,2,\ldots,T$, by

$$\mathbf{y}_i(t) = \begin{cases} (1-\psi_{i-1})\mathbf{y}_{i-1}(t-1) - \mathbf{z}_{i-1}(t-1) & i = 2,3,\ldots,n; \\ (1-\psi_0)\sum_{j=1}^n \alpha_j \mathbf{y}_j(t-1) - \mathbf{z}_0(t-1) & i = 1. \end{cases} \tag{1}$$

The second equation describes the dynamics of lambs that become age-1 ewes in the following year. Also, ewes of age n are sold, so that for each $t = 0,1,\ldots,T-1$, we have

$$(1-\psi_n)\mathbf{y}_n(t) - \mathbf{z}_n(t) = 0.$$

Figure EC.1 provides a state-transition diagram to aid in understanding.

The parameter ℓ^* represents the target flock size — in terms of the number of ewes — to manage vegetation growth on the solar farm. The flock size cannot exceed this value, i.e.,

$$\sum_{i=1}^n \mathbf{y}_i(t) \le \ell^* \qquad \forall t = 0,1,\ldots,T-1.$$

The vector giving the initial flock census is given by $\mathbf{y}(0) = \ell + \mathbf{x}$, where $\ell \in [0,\infty)^n$ is a vector giving the initial number of ewes on hand, broken down by age $i = 1,2,\ldots,n$, and $\mathbf{x}_i$ gives the number of age-i ewes purchased at time 0 at a cost of c_i per head, $i = 1,2,\ldots,n$. The vector $\ell = 0$ if no ewes are initially on hand.

The objective is to maximize the total discounted (at rate r) profit,

$$-\sum_{i=1}^n c_i \mathbf{x}_i + \sum_{t=0}^{T-1} (1+r)^{-(t+1)} \left[\sum_{i=0}^n p_i \mathbf{z}_i(t) + p^R(1-\psi_0) \sum_{i=1}^n \alpha_i \mathbf{y}_i(t) - c^H \sum_{i=1}^n \mathbf{y}_i(t) \right].$$

Here, we account for overwintering and annual costs based on the flock composition at the start of each year. The complete linear programming (LP) formulation is thus to

$$
\begin{array}{lllll}
\max & -\sum_{i=1}^{n} c_i \mathbf{x}_i + \sum_{t=0}^{T-1} (1+r)^{-(t+1)} \left[\sum_{i=0}^{n} p_i \mathbf{z}_i(t) + p^R (1-\psi_0) \sum_{i=1}^{n} \alpha_i \mathbf{y}_i(t) - c^H \sum_{i=1}^{n} \mathbf{y}_i(t) \right] & & & \\
\text{s/t} & (1-\psi_0) \sum_{j=1}^{n} \alpha_j \mathbf{y}_j(t-1) - \mathbf{z}_0(t-1) - \mathbf{y}_1(t) & = 0 & & t = 1, 2, \ldots, T \\
& (1-\psi_{i-1}) \mathbf{y}_{i-1}(t-1) - \mathbf{z}_{i-1}(t-1) - \mathbf{y}_i(t) & = 0 & & t = 1, 2, \ldots, T; i = 2, 3, \ldots, n \\
& (1-\psi_n) \mathbf{y}_n(t) - \mathbf{z}_n(t) & = 0 & & t = 0, 1, \ldots, T-1 \\
& \sum_{i=1}^{n} \mathbf{y}_i(t) & \le \ell^* & & t = 0, 1, \ldots, T-1 \\
& \mathbf{y}_i(0) & = \ell_i + \mathbf{x}_i & & i = 1, 2, \ldots, n \\
& \mathbf{x}_i \ge 0 & & & i = 1, 2, \ldots, n \\
& \mathbf{y}_i(t) \ge 0 & & & t = 0, 1, \ldots, T-1; i = 1, 2, \ldots, n \\
& \mathbf{z}_i(t) \ge 0 & & & t = 0, 1, \ldots, T-1; i = 0, 1, \ldots, n.
\end{array}
\tag{2}
$$

If a ewe does not have sufficient lambs over the course of its life, then it is not feasible to maintain a solar-grazing operation. In this case, assuming it is still profitable to raise lambs, numerical results show that the optimal solution to Problem (2) is simply to sell any existing ewes once they reach their maximum age of n years, while selling all lambs. To avoid this trivial solution, we impose a condition, satisfied for reasonable choices of parameters, that ewes have sufficient offspring. To that end, let $\zeta_0 \triangleq 1$ and $\zeta_i \triangleq (1-\psi_1)(1-\psi_2)\cdots(1-\psi_i)$ for $i = 1, 2, \ldots, n$, so that ζ_i represents the probability that a ewe survives at least until the end of its ith year. Define

$$
R_n \triangleq 1 + (1-\psi_1) + (1-\psi_1)(1-\psi_2) + \cdots + (1-\psi_1)(1-\psi_2)\ldots(1-\psi_{n-1}) = \sum_{i=0}^{n-1} \zeta_i,
$$

which is the expected lifetime of a ewe in the absence of being sold before its maximal lifetime, since it is a sum of tail probabilities. The expected number of ewe lambs surviving their first year that are born to a single ewe over the course of the ewe's lifetime is $(1-\psi_0)\sum_{j=1}^{n} \alpha_j \zeta_{j-1}$. For ewes to be able to replace themselves, this quantity must be at least 1. In the case of constant reproduction rate α, this perspective establishes a minimum reproduction rate, $\alpha \ge [(1-\psi_0)\sum_{j=1}^{n} \zeta_{j-1}]^{-1}$, for which the endeavor remains possible for the farmer.

Assumption 1 **(Flock Stability)**. *Ewes replace themselves, i.e.,* $(1-\psi_0)\sum_{j=1}^{n} \alpha_j \zeta_{j-1} \ge 1$.

Assumption 1 is not sufficient to ensure that the solar-grazing operation is profitable. Indeed, if the selling price vector p is too small, then it costs more to over-winter ewes, so it is more cost

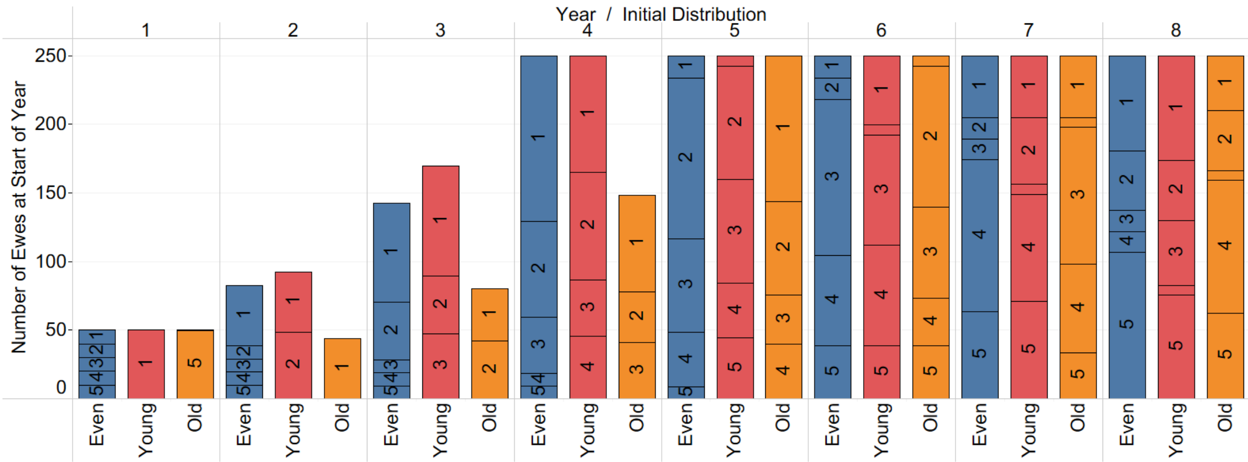

Figure 1 **Ewe population by age and year for an *even* initial ewe distribution $\mathbf{y}(0) = (10, 10, 10, 10, 10)$, *young* initial distribution $\mathbf{y}(0) = (50, 0, 0, 0, 0)$, and *old* initial distribution $\mathbf{y}(0) = (0, 0, 0, 0, 50)$ under the maximum-lifetime policy. Note: ewe lambs, which are less than one year old, are not shown, and age labels are omitted when cohorts are very small.**

effective to simply sell all ewes immediately. Moreover, startup costs and annual fixed costs must also be considered. We explore an example that takes all these factors into account in Section 3.2.

The following results are for a representative case (see §3.2 for discussion of these choices) with the parameters $n = 5$ years (which reflects a maximum age of 6 years, since we begin indexing at 0), birth rate per year of ewe lambs $\alpha_i = 0.9$, procurement cost $c_i = \$389$, selling price $p_i = \$100$ $(i = 1, 2, \ldots, 5)$, loss rate $\psi_i = 0.03$ $(i = 0, 1, \ldots, 5)$, lamb selling price $p_0 = p^R = \$250$, annual per-head cost $c^H = \$66$, discount rate $r = 0.05$, time horizon $T = 20$ years, and maximum flock size $\ell^* = 250$.

Given an initial flock population $\ell = (10, 10, 10, 10, 10)$, the optimal solution with these parameters is to purchase 200 year-1 ewes to immediately attain the desired flock size $\ell^* = 250$, then to adopt a *Maximum-Lifetime* policy, in that ewes are only sold when they reach the maximum possible age, and excess ewe lambs beyond those needed to maintain the population of ℓ^* are sold at the end of their first year.

Suppose that we want to explore the potential to simply "breed our own flock," as opposed to buying enough ewes from the outset to reach the desired flock size ℓ^*. To that end, we constrain the initial purchase variables, $x_i = 0$ for all i. The solution is depicted in Figure 1, which illustrates how the flock size and age distribution evolve over time for different initial age distributions. The results show that the initial distribution primarily affects the speed at which the flock approaches its maximum size: starting from a flock consisting solely of age-5 ewes delays reaching the maximum size by one year relative to starting from age-1 ewes or from an evenly distributed age structure. Moreover, the flock size increases geometrically at first, before reaching a plateau at the desired flock size of 250.

3.1. The Initial Purchase Decision

Given this model of flock dynamics, we can provide some structure on the initial flock procurement decision. For any fixed initial procurement vector $\mathbf{x}$, let $g(\mathbf{x})$ denote the optimal objective function value of the LP Formulation (2) where $\mathbf{x}$ is held fixed.

PROPOSITION 1. *Given a fixed initial flock size ℓ, the optimal profit of Problem* (2)*, $g(\mathbf{x})$, is a concave and piecewise affine function of the initial procurement vector* $\mathbf{x}$.

Proposition 1 establishes a "diminishing returns" property for initial flock purchases. The result assumes that all decisions are made optimally with respect to the LP formulation, but the observation applies much more generally. For example, suppose we restrict attention to the maximum-lifetime policy in which sales are allowed only for ewe lambs and for ewes in their final year n. Accordingly, we restrict the feasible set by imposing the additional constraints that $\mathbf{z}_1(t) = \mathbf{z}_2(t) = \cdots = \mathbf{z}_{n-1}(t) = 0$ for all t. Augmenting the LP formulation (2) with these linear constraints yields another LP. It follows immediately that the optimal profit remains a concave, piecewise-affine function of $\mathbf{x}$. We omit the proof.

COROLLARY 1. *Fix the initial flock size ℓ. Suppose we use the maximum-lifetime policy. Then, the optimal profit of Problem* (2) *is a concave and piecewise affine function of* $\mathbf{x}$.

Along the same lines, one can extend Corollary 1 to a broad class of policies beyond the maximum-lifetime policy, as long as the resulting formulation remains a linear program. For example, this holds for any policy under which $\mathbf{z}_i(t)$ can be expressed as a linear function of the $\mathbf{y}$-variables for all i and t.

The diminishing returns property, i.e., concavity, is perhaps most salient in a setting where the initial procurement decision is restricted to be one-dimensional. Indeed, suppose that a farmer will decide the *number* of ewes to purchase, say x, while the proportion of ewes of each age $(\gamma_1, \gamma_2, \ldots, \gamma_n)$ is fixed, a situation that can arise when buying ewes at auction. The initial procurement is then restricted to be of the form $\mathbf{x} = x(\gamma_1, \gamma_2, \ldots, \gamma_n)$. The function $\tilde{g}(x)$ that maps the one-dimensional decision variable x to the optimal profit is simply the restriction of the function $g(\mathbf{x})$ to a line, so that if g is concave then so is $\tilde{g}$.

COROLLARY 2. *Fix the initial flock size ℓ. Suppose the initial procurement of ewes is given by $x(\gamma_1, \gamma_2, \ldots, \gamma_n)$ for fixed values $\gamma_1, \gamma_2, \ldots, \gamma_n$. Then, the optimal profit of the Problem* (2) *is a concave and piecewise affine function of x.*

Corollary 2 makes clear that the marginal change in profit is non-increasing in the initial procurement size, x; i.e., the value of each additional ewe is non-increasing. The numerical results in Section 3.2 showcase this relationship.

3.2. Numerical Illustration

To illustrate the impact of our model on the initial flock procurement decision, we present a use case based on our research partner — who both operates a solar grazing connection platform and grazes his own sheep across several sites — for a hypothetical $A = 100$-acre site, using parameters representative of his aggregated operations. Our estimates align with those proposed by Kochendoerfer and Thonney (2021) and Shiflett (2022), who provide additional practical insights into solar grazing. The site's grazing capacity is 2.5 breeding ewes per acre, counting ewes together with their offspring and the rams required for breeding (approximately 2–3 per 100 ewes). Accordingly, the site can support a maximum of $\ell^* = 250$ ewes. The farmer starts with a flock of $\ell_* = 50$ ewes. We assume that the initial flock and any additional ewes acquired at the start of the horizon are evenly distributed over ages 1 through 5, so that $\gamma_1 = \gamma_2 = \cdots = \gamma_5 = 0.2$ in the context of Corollary 2.

Each ewe produces $\alpha = 0.9$ ewe lambs and $\alpha = 0.9$ ram lambs annually, yielding a total reproduction rate of $2\alpha = 1.8$ lambs per ewe per year. Each year, $\psi = 3\%$ of both lambs and ewes are lost due to predation or disease. While the model presented in Section 3 allows these parameters to vary by ewe age, we assume they are age-invariant for simplicity. Ewes in their sixth year are sold before winter. Following the preferences of our farm partner, we assess undiscounted cash flows (i.e., $r = 0$). The qualitative insights remain consistent under a positive discount rate.

Besides the cost of procuring sheep, the farmer incurs three additional categories of expenses: upfront costs, annual fixed costs, and annual per-head costs. Although the model introduced at the beginning of Section 3 does not require upfront or annual fixed costs for its structural analysis, we include them here to evaluate the overall economics of a solar grazing operation. The cost parameters are as follows:

Upfront costs $c^U = \$28,500$. Most of this cost is for sheep handling equipment.

Fixed annual net cost $c^Y = \$52,000$. To estimate this amount, the farmer incurs a cost of \$40,000 — i.e., 100 acres at \$400 per acre — for vegetation management without any sheep. In addition, the farmer faces significant operational expenses totaling approximately \$60,000, including insurance, water hauling, staffing or labor, and help from a herding dog. The majority of this expense is attributed to the base level of labor required to operate the solar grazing system, which can be performed either by hired workers or by the farmer and their family. Our farm partner estimated that managing a 250-sheep solar grazing operation requires roughly 50% of a full-time equivalent position during the winter months and 25% during the summer. For a 100-acre site, this includes the equivalent of roughly 10 working days annually devoted to site maintenance tasks other than mowing. These costs are partially offset by $\$48,000$ in maintenance contract revenue, which includes \$38,000 for vegetation management and \$10,000 for ancillary services such as light clearing work outside the fenced area.

Per head annual costs $c^H = \$66$. Our farm partner estimates that maintenance costs vary linearly with the number of sheep and that certain site maintenance would still be required even with a full flock. After accounting for this additional work, each of the 250 ewes saves \$144 in maintenance costs. The farmer also incurs several per-head annual expenses, including transportation, veterinary care, mineral supplements, overwintering, and labor totaling \$210 per ewe. Among these, overwintering is the largest expense, covering the cost of hay consumed when the sheep are not grazing in the field.

Procurement cost $c = \$389$. A ewe costs approximately \$300 for the farmer to procure through a livestock auction or private sale. Livestock purchases are financed through a bank loan at an 8% interest rate, amortized over six years, resulting in a total payment of \$389 based on the standard amortization formula.

Because our farmer partner has other entrepreneurial commitments, much of the labor must be performed by hired farm workers or family members. The operation's profitability appears much more favorable for farmers who perform the work themselves, possibly including vegetative site maintenance beyond sheep grazing and producing hay for overwintering as part of their broader farm operations.

Additionally, the farmer earns revenue from the following sources:

Proceeds from a lamb $p^E = p^R = \$250$. Selling a lamb generates \$250. Price discrepancies between the purchase and sale of ewes (i.e., \$389 vs. \$250) may stem from factors such as transportation costs, auction fees, interest on loans, or the cost of breeding quality certifications.

Proceeds from an adult sheep $p^S = \$100$. An older ewe sells for approximately \$100.

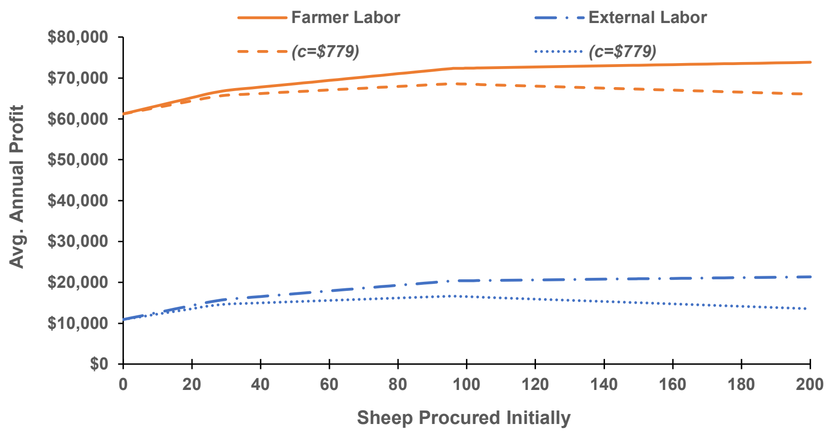

Figure 2 Average annual profit over the first ten years as a function of the initial flock procurement quantity, shown separately for scenarios in which labor is performed by the farmer or outsourced to external workers.

Figure 2 illustrates the piecewise linear profit function as described in Corollary 2 for this problem under two scenarios: when the annual labor expense — which includes both fixed yearly and per-head components — is considered as either an external cost or part of the farmer's earnings. In both scenarios, we observe significant benefits from purchasing 200 sheep immediately. In the scenario with external labor costs and $c = \$389$, the operation's average annual profit more than doubles when sheep are procured upfront, compared to allowing the flock to grow solely through reproduction. These gains come from reduced mowing costs — estimated at $144 per ewe per year — and from additional lambs available for sale. Without procuring any additional sheep, the flock would reach its maximum size only at the end of Year 3.

While this particular set of parameters guides the farmer towards procuring enough ewes to reach the maximum flock size before beginning operations, other parameters may lead to an optimal solution in which the farmer purchases only part of the flock and relies on reproduction for growth. For example, if the procurement cost doubles (i.e., $c = \$779$) — perhaps due to interest in higher-quality breeding stock or a specific heritage breed — the optimal solution becomes purchasing only 99 ewes. The flock reaches its maximum size of 250 at the end of year 1. This strategy improves the average annual profit by approximately $2,833 ($2,306 with labor performed by the farmer) compared to the alternative of initially purchasing 200 sheep.

3.3. Implementation

This research is motivated by a collaboration with United Agrivoltaics, a start-up solar grazing platform that also serves as a key setting for implementing our insights. In addition to connecting solar developers with sheep farmers — some of whom already have flocks and others who are just starting out — United Agrivoltaics helps farmers to navigate legal contracts, training and compliance, and insurance requirements. As of 2025, the United Agrivoltaics network includes approximately 200 farmers distributed across the United States and also offers consulting services related to solar grazing. In this section, we describe how our research has created impact through United Agrivoltaics, including the deployment of a spreadsheet model for procurement decisions, consulting support for utility-scale operations, and assistance with site approval processes.

3.3.1. Spreadsheet Model for Procurement Our early research collaboration with United Agrivoltaics involved documenting the operational details of solar grazing in spreadsheet models using either the state dynamics of the LP formulation (2) or an approximation that does not track ewes by age. United Agrivoltaics quickly recognized an opportunity to adapt this model for onboarding new farm partners. In response, we developed a version of the spreadsheet that became part of their standard training materials. Today, over 70 farm partners with active operations have access to the spreadsheet. United Agrivoltaics estimates that at least half of them have actively used the model to guide their initial flock procurement decisions and forecast operational cash flows.

The experience of one farmer in Georgia illustrates the value of our model. This young man hoped to begin solar grazing but had previously been denied a loan, as the bank — despite being open to financing other types of livestock — was averse to financing sheep, calling them "unbankable." However, with the support of our model, which documented the economics of solar grazing, the bank approved a loan to procure a full flock. The farmer began operations on six solar sites and has since expanded to 12, now working as a full-time solar grazier. He later recommended the spreadsheet to another farmer in Georgia, who similarly used the model to secure a loan and launch their own solar grazing operation.

3.3.2. Consulting for Utility-Scale Operators United Agrivoltaics provides consulting services to utility-scale clients, including one Midwestern utility that has added over 600 MW of solar energy generation capacity in recent years. In an effort to achieve both cost savings and environmental benefits, the utility has explored the use of solar grazing. United Agrivoltaics has used our spreadsheet model of the farmer's procurement decision to collaborate with the utility in designing contracts that are economically beneficial for farmers, while also taking advantage of potential tax savings under a structure in which the utility retains ownership of the sheep.

In a consulting engagement with another utility-scale operation, a United Agrivoltaics consultant used our model to assess a specific solar grazing contract and identify its issues. In this case, the asset owner — seeking to promote solar grazing — had contracted a third-party vegetation management company, which negotiated aggressively with a farmer, resulting in an unsustainably low maintenance fee per acre. Conversations with the farmer revealed that they had dramatically underestimated maintenance costs during the flock growth phase and were at risk of bankruptcy within the first few years. Consequently, the utility revisited the arrangement with both the third-party company and the farmer to make the contract economically sustainable.

3.3.3. Consulting for Site Approval United Agrivoltaics also provides consulting services to solar developers navigating site permitting and variance approval processes at the local level — procedures that are often protracted and contentious. Developers often face strong opposition due to concerns about farmland being diverted from agricultural use. In several communities across the United States, solar projects have initially been rejected or subjected to moratoriums, only to be approved later when agricultural activities like solar grazing were incorporated. Through Agriculture Impact Statements and related analyses, United Agrivoltaics helps demonstrate that solar grazing can provide greater agricultural value than traditional crops such as corn or soybeans, thereby reducing legal and political barriers to project approval.

For example, a proposal for a 46-acre solar site in the town of Lockport, NY, sparked intense local protests that led to a moratorium on solar development. After two more years of contentious public meetings, a consultant used outputs from our model to demonstrate the agricultural value of solar grazing, ultimately securing approval from the town board. The solar site was completed in 2024, and sheep began grazing it in 2025. United Agrivoltaics has also used our model to support the approval processes for several other sites across the Mid-Atlantic and Northeastern United States, many of which are still under development.

4. Steady-State Flock Dynamics

In the time-dependent model of Section 3, we adopted the maximum-lifetime policy both in Corollary 1 and in the numerical illustrations shown in Figure 1, without formally evaluating its performance. We now provide strong evidence for the effectiveness of this policy in a steady-state flock management setting where both sales and retention decisions are stationary in time. This restriction leads to a steady-state linear program that captures long-run flock maintenance decisions.

We first show that this simplification is without loss of optimality: any objective value achievable through a complex, time-varying, multi-year cycle can also be achieved by a steady-state policy. Thus, in steady state, there is no need to pursue complicated policies to maximize profits. We then prove the maximum-lifetime policy is optimal under mild regularity conditions, which is the main focus of the section. Finally, we complement our theoretical results with numerical illustrations.

4.1. Analytical Results

Our focus is on steady-state sales decisions, so we first introduce the following LP formulation, which is a natural analogue of Problem (2) over a cycle of length T years:

$$
\begin{array}{lllll}
\max & \frac{1}{T}\sum_{t=0}^{T-1}\left[\sum_{j=0}^{n} p_j \mathbf{z}_j(t) + p^R(\mathbf{z}_0(t)+\mathbf{y}_1(t))\right] & & & \\
\text{s/t} & (1-\psi_0)\sum_{j=1}^{n}\alpha_j \mathbf{y}_j(t-1) - \mathbf{z}_0(t-1) & = \mathbf{y}_1(t) & t=1,2,\ldots,T & \\
 & (1-\psi_{i-1})\mathbf{y}_{i-1}(t-1) - \mathbf{z}_{i-1}(t-1) & = \mathbf{y}_i(t) & t=1,2,\ldots,T; i=2,3,\ldots,n & \\
 & (1-\psi_n)\mathbf{y}_n(t) - \mathbf{z}_n(t) & = 0 & t=0,1,\ldots,T-1 & \\
 & \sum_{i=1}^{n}\mathbf{y}_i(t) & = \ell^* & t=0,1,\ldots,T-1 & \\
 & \mathbf{y}_i(T)-\mathbf{y}_i(0) & = 0 & i=1,2,\ldots,n & \\
 & \mathbf{y}_i(t)\geq 0 & & t=0,1,\ldots,T; i=1,2,\ldots,n & \\
 & \mathbf{z}_i(t)\geq 0 & & t=0,1,\ldots,T-1; i=0,1,\ldots,n. &
\end{array} \tag{3}
$$

Compared with the formulation (2), (3) differs in several key aspects:

- The objective is expressed in terms of average profit rather than discounted profit.
- We rule out initial purchases by setting $\mathbf{x}=\mathbf{0}$.
- We impose a cyclic nature to flock levels, $\mathbf{y}_i(T)=\mathbf{y}_i(0)$ for all $i=1,2,\ldots,n$, ensuring that the length-T trajectory repeats itself.
- The flock-size constraint is imposed as a binding constraint.
- The flock size is fixed over time, so the term $c^H\sum_{i=1}^{n}\mathbf{y}_i(t)$ is a constant that does not affect the optimization and can therefore be omitted from the objective function. In addition, the term $p^R(1-\psi_0)\sum_{j=1}^{n}\alpha_j\mathbf{y}_j(t)$ in the objective simplifies to $p^R(\mathbf{z}_0(t)+\mathbf{y}_1(t))$, using (1) and the cyclic constraints $\mathbf{y}_i(T)=\mathbf{y}_i(0)$.

We next show that Problem (3) is equivalent to the following steady-state formulation, which restricts the decision variables to be time-invariant (i.e., independent of t):

$$\begin{array}{llll} \max & \sum_{j=0}^{n} p_j z_j + p^R(z_0+y_1) & & \\ \text{s/t} & (1-\psi_0)\sum_{j=1}^{n}\alpha_j y_j - z_0 & = y_1 & \\ & (1-\psi_{i-1})y_{i-1} - z_{i-1} & = y_i & i=2,3,\ldots,n \\ & (1-\psi_n)y_n - z_n & = 0 & \\ & \sum_{i=1}^{n} y_i & = \ell^* & \\ & y_i \geq 0 & & i=1,2,\ldots,n \\ & z_i \geq 0 & & i=0,1,\ldots,n. \end{array} \tag{4}$$

Proposition 2. *The optimal objective values of the time-dependent formulation* (3) *and the steady-state formulation* (4) *are identical.*

Proposition 2 shows that we can restrict our attention to steady-state solutions without loss of optimality. These solutions are also simpler to implement than multi-year cycles, which require managing age profiles that vary from year to year.

We next characterize the optimal solution to the steady-state formulation (4) under certain assumptions on the input parameters. Under certain conditions, we prove that the optimal solution to the steady-state formulation (4) follows a maximum-lifetime structure, in which only lambs and ewes in the final year n are subject to sales (i.e., $z_1^* = z_2^* = \cdots = z_{n-1}^* = 0$). Throughout the rest of the paper, we refer to this as the *maximum-lifetime* policy. Interestingly, the maximum-lifetime policy with $z_0 > 0$ (i.e., selling lambs) appears to be unique to our problem, which incorporates a constraint on the flock size.

Theorem 1. *Suppose $(1-\psi_0)\sum_{j=1}^{n}\alpha_j\zeta_{j-1} - 1 \geq 0$ and that the prices $\{p_i\}_{i=0}^{n}$ are decreasing and convex with respect to (w.r.t.) i. In addition, suppose the reproduction rates $\{\alpha_i\}_{i=1}^{n}$ are non-decreasing in i and the loss rates $\{\psi_i\}_{i=1}^{n}$ are non-increasing in i. Then, there exists an optimal solution where $z_1^* = z_2^* = \cdots = z_{n-1}^* = 0$. Thus, $y_i^* = (1-\psi_{i-1})y_{i-1}^*$ for $i = 2,3,\ldots,n$ which, with the constraint on flock size that $\sum_{i=1}^{n} y_i = \ell^*$ determines the age distribution of sheep. Moreover,*

$$z_0^* = \left((1-\psi_0)\sum_{j=1}^{n}\alpha_j\zeta_{j-1} - 1\right) R_n^{-1}\ell^*, \qquad z_n^* = \zeta_n R_n^{-1}\ell^*.$$

In the special case where all α_i and ψ_i are homogeneous, i.e., $\alpha_i = \alpha$ and $\psi_i = \psi$ for all i,

$$z_0^* = \left[(1-\psi)\alpha - (R_n)^{-1}\right]\ell^*, \quad z_n^* = \frac{(1-\psi)^n \ell^*}{R_n}.$$

We will discuss the relevance of Theorem 1's remaining assumptions on $\{p_i\}_{i=0}^n$, $\{\alpha_i\}_{i=1}^n$ and $\{\psi_i\}_{i=1}^n$ in Section 4.2. The necessity of these conditions becomes evident in the proof of Theorem 1, provided in the appendix. A key observation is that (4) can be reformulated purely in terms of the variables y_i. The conditions on $\{p_i\}_{i=0}^n$, $\{\alpha_i\}_{i=1}^n$ and $\{\psi_i\}_{i=1}^n$ are essential to ensure a certain monotonicity in the objective function of the reformulated problem. If any of these assumptions on the structure of input parameters are violated, the conclusion of Theorem 1 may no longer hold. We illustrate this with two counterexamples.

Example 1. Suppose $n=2$, $\psi_0=\psi_1=\psi_2=0$, and $p^R=0$. The following two solutions are feasible:

$$z_0=(\alpha_1+\alpha_2-1)\frac{\ell^*}{2}-(1+\alpha_2-\alpha_1)\frac{\epsilon}{2},\quad z_1=\epsilon,\quad z_2=y_2=\frac{\ell^*}{2}-\frac{\epsilon}{2},\quad y_1=\frac{\ell^*}{2}+\frac{\epsilon}{2};$$
$$z_0'=(\alpha_1+\alpha_2-1)\frac{\ell^*}{2},\quad z_1'=0,\quad z_2'=y_1'=y_2'=\frac{\ell^*}{2},$$

(assuming $\alpha_1+\alpha_2>1$ and ϵ is sufficiently small, to ensure all y_i, z_i are non-negative). It follows that

$$\sum_{j=0}^{2} p_j z_j' - \sum_{j=0}^{2} p_j z_j = \left[p_0\,(1+\alpha_2-\alpha_1)+p_2-2p_1\right]\frac{\epsilon}{2}, \tag{5}$$

which can be either positive or negative, depending on the relation between α_1 and α_2. For instance, when $p_0=p_1=p_2$, the sign of (5) equals the sign of $\alpha_2-\alpha_1$. In particular, when $\alpha_2<\alpha_1$, the conclusion of Theorem 1 no longer holds. This is intuitive because, for instance, when $\alpha_2=0$ (i.e., when the reproduction rate for ewes in their final year is zero), it becomes more beneficial to sell more ewes a year earlier to reduce the number of ewes in their final year.

Example 2. Suppose $n=2$, $\psi_0=0$, $\alpha_1=\alpha_2=1$, and $p^R=0$. The following two solutions are feasible:

$$z_0=\frac{1-\psi_1}{2-\psi_1}\ell^*-\frac{1}{2-\psi_1}\epsilon,\quad z_1=\epsilon,\quad z_2=\frac{(1-\psi_1)(1-\psi_2)}{2-\psi_1}\ell^*-\frac{1-\psi_2}{2-\psi_1}\epsilon,$$
$$y_1=\frac{1}{2-\psi_1}\ell^*+\frac{1}{2-\psi_1}\epsilon,\quad y_2=\frac{1-\psi_1}{2-\psi_1}\ell^*-\frac{1}{2-\psi_1}\epsilon;$$
$$z_0'=\frac{1-\psi_1}{2-\psi_1}\ell^*,\quad z_1'=0,\quad z_2'=\frac{(1-\psi_1)(1-\psi_2)}{2-\psi_1}\ell^*,\quad y_1'=\frac{1}{2-\psi_1}\ell^*,\quad y_2'=\frac{1-\psi_1}{2-\psi_1}\ell^*.$$

It follows that

$$\sum_{j=0}^{2} p_j z_j' - \sum_{j=0}^{2} p_j z_j = \left[\frac{1}{2-\psi_1}p_0+\frac{1-\psi_2}{2-\psi_1}p_2-p_1\right]\epsilon, \tag{6}$$

which can be either positive or negative, depending on the relation between ψ_1 and ψ_2. For instance, when $p_0=p_1=p_2$, the sign of (6) equals the sign of $\psi_1-\psi_2$. In particular, when $\psi_1<\psi_2$, the conclusion of Theorem 1 no longer holds. This is also intuitive because, for instance, when $\psi_2 = 1$ (i.e., a 100% loss of ewes in their final year), it becomes more beneficial to sell more ewes a year earlier to reduce the number of ewes in their final year.

4.2. Numerical Illustration

In this section, we illustrate the practical relevance of the flock maintenance model and evaluate the performance of the maximum-lifetime policy. We begin by discussing the relevance of the three key assumptions in ensuring the optimality of the maximum-lifetime policy in Theorem 1: (1) $\{\alpha_j\}$ is increasing in j, (2) $\{\psi_j\}$ is decreasing in j, and (3) $\{p_j\}$ is decreasing and convex in j.

First, the increasing trend in birth rate $\{\alpha_j\}$ with age is practical. As noted in Notter (2000), ewe prolificacy for Targhee sheep peaks between ages 4 and 8. Compared to mature ewes, 1-year-old Targhee ewes produce approximately 0.6 fewer lambs, 2-year-olds about 0.3 fewer, and 3-year-olds roughly 0.1 fewer. This pattern reflects physiological maturity, as younger ewes prioritize growth over reproduction, gradually shifting energy allocation to reproductive functions as they mature.

Second, the decline in loss rate $\{\psi_j\}$ with age is both reasonable and empirically observed, e.g., Flay et al. (2021) report decreasing annual mortality from age 1 to age 6. Biologically, younger ewes are still developing, and early breeding can be stressful, increasing the risk of mortality. In contrast, mature ewes experience fewer lambing complications, contributing to lower loss rates.

Third, the assumption that price decreases in a convex manner with age is well-supported. There is a significant price gap between lamb (under 12 months) and mutton (over two years). Mutton is typically much cheaper than lamb, with recent data showing mutton prices 76% lower than lamb prices (FarmingAhead 2024). For example, the price sequence $p_0 = \$250$ and $p_1 = p_2 = \cdots = p_n = \$100$ reflects a realistic pricing structure. This price disparity arises largely because lamb is prized for its tender texture and milder flavor, and its supply is naturally limited—each sheep produces only a limited number of lambs annually. In contrast, mutton production can be ramped up more rapidly, and holding sheep beyond the lamb stage is less cost-effective due to slowed growth and increased maintenance costs. This creates a convex age-price curve: prices drop sharply after the lamb phase before stabilizing near a low salvage value. Even if quality differences are subtle, consumer perception and labeling drive a disproportionate price drop as the animal moves into a less desirable market class.

Our assumptions are realistic, but one can imagine modest deviations from the assumed properties that could render the maximum-lifetime (ML) policy sub-optimal. Here, we explore the degree of this sub-optimality through numerical examples, constructed from conceivable parameter values.

We also aim to identify conditions under which the ML policy may perform poorly. Unless otherwise specified, we fix $n=5$, $p^R=250$, and $\psi_0=0.1$, while varying $\{p_j\}_{j=0}^n$, $\{\alpha_j\}_{j=1}^n$, and $\{\psi_j\}_{j=1}^n$.

- **Scenario 1**: We consider the most plausible parameter values in this scenario:

$$(p_0,\dots,p_n)=(250,100,100,100,100,100),$$
$$(\psi_1,\dots,\psi_n)=(0.03,0.03,0.03,0.03,0.03),$$
$$(\alpha_1,\dots,\alpha_{n-1})=(0.6,0.9,1.1,1.2),$$

  and we vary the birth rate in the final year, α_n, from 1.1 to 0.

- **Scenario 2**: This scenario is similar to Scenario 1, except we consider keeping sheep for one additional year, i.e., $n=6$. In this case, birth rates rise to a peak and then decline in the final two years, and the price of a ewe decreases from \$100 to \$75 in the final year:

$$(p_0,\dots,p_n)=(250,100,100,100,100,100,75),$$
$$(\psi_1,\dots,\psi_n)=(0.03,0.03,0.03,0.03,0.03,0.03),$$
$$(\alpha_1,\dots,\alpha_{n-1})=(0.6,0.9,1.1,1.2,1.1),$$

  and we vary the birth rate in the final year, α_n, from 1.1 to 0.

- **Scenario 3**: This scenario is similar to Scenario 1, except we stress-test the final year loss rate instead of the birth rate:

$$(p_0,\dots,p_n)=(250,100,100,100,100,100),$$
$$(\alpha_1,\dots,\alpha_n)=(0.6,0.9,1.1,1.2,0.6),$$
$$(\psi_1,\dots,\psi_{n-1})=(0.03,0.03,0.03,0.03),$$

and we vary the loss rate in the final year, ψ_n, from 0.03 to 1. Note that to observe any effect, we must choose a low value for α_n, such as 0.6 (compared to 1.2); otherwise, the performance ratio remains close to one.

We report the performance ratios of the ML policy relative to the optimal LP value in Table 1. Overall, the ML policy continues to perform well even when the assumptions are violated under realistic parameter values, and it can perform very close to optimality when the deviations are small. These findings are both timely and relevant, as the ML policy remains a widely used guideline in practice. This is largely because ewes are most productive between ages 3 and 6 and experience minimal loss rates during this period, as discussed earlier. Consequently, routine culling or sale by age 6 is standard to maintain flock productivity (e.g., OSU Sheep Team 2018). Both our theoretical analysis and numerical results support the strong performance of this policy.

Table 1 Numerical results for Scenarios 1-3.

α_n	1.1	1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0
Ratio (**Scenario 1**)	1	1	1	0.99	0.97	0.94	0.92	0.90	0.88	0.85	0.83	0.81
Ratio (**Scenario 2**)	1	1	0.99	0.97	0.96	0.94	0.92	0.90	0.89	0.87	0.85	0.84
ψ_n	0.03	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Ratio (**Scenario 3**)	0.94	0.94	0.94	0.94	0.93	0.93	0.92	0.92	0.91	0.91	0.90	0.90

Given the robust performance of the ML policy for realistic parameters, one might ask how poorly it might perform for other parameters, e.g., in case a similar model is employed for livestock other than sheep. To that end, we derive a worst-case performance ratio bound for the ML policy that is independent of the price parameters. In the present setting of sheep flocks, this bound tends to be tight when the optimal LP solution retains only first-year ewes (i.e., $y_j^* = 0$ for all $j > 1$), deviating significantly from the structure of the ML policy suggested in Theorem 1.

PROPOSITION 3. *Assume that $p^R = p_0 \geq p_i$ for all $i = 1\ldots, n$. The performance ratio of the ML policy value relative to the optimal LP value is at least*

$$\frac{\left(2(1-\psi_0)\sum_{j=1}^{n}\alpha_j\zeta_{j-1} - 1\right) R_n^{-1}}{\max\limits_{j=1,\ldots,n}\{2(1-\psi_0)\alpha_j - \psi_j\}}. \tag{7}$$

The denominator of (7) provides an upper bound on a scalar multiple of the optimal LP value. The intuition behind this upper bound is as follows: we start with the full breeding population (ℓ^*), assume that all sheep are sold at the highest possible price p_0, and take the most optimistic values for reproduction and loss — specifically, the highest reproduction rate $\max_{j=1}^n \alpha_j$ and the smallest loss rate $\min_{j=1}^n \psi_j$, which yields the largest population gain rate $\max_{j=1,\ldots,n}\{2(1-\psi_0)\alpha_j - \psi_j\}$.

The bound is tight when (1) the optimal LP solution retains only first-year ewes; (2) $p_0 = p_1$; and (3) first-year ewes exhibit both the highest reproduction rate and lowest loss rate across all age groups. This aligns with the intuition discussed above.

One example of such a parameter setting arises when $p_0 = p_1 = \ldots = p_{n-1} >> p_n$, with $\{\alpha_j\}_{j=1}^n$ decreasing in j and $\{\psi_j\}_{j=1}^n$ increasing in j. In this case, the price structure is substantially altered, and the monotonicity of both $\{\alpha_j\}$ and $\{\psi_j\}$ is the exact opposite of the assumptions that guarantee the optimality of the ML policy. Specifically, when $p_0 = p_1 = \ldots = p_{n-1} = 250$, $p_n = 0$, $\alpha = (1.4, 1.2, 1.0, 0.8, 0.6)$ and $\psi = (0.01, 0.02, 0.03, 0.04, 0.05)$, the performance ratio is 0.6414, exactly matching the lower bound established in Proposition 3 .

5. Stochastic Flock Dynamics

The models analyzed thus far rely on fluid dynamics, replacing stochastic births and losses with deterministic, continuous-valued dynamics. We now extend those models to explicitly incorporate stochastic, integer-valued births and losses. This extension allows us to assess whether our earlier conclusions remain valid in a more realistic and granular setting.

Our primary finding is that the stochastic model aligns with the steady-state LP in (4). In particular, the maximum achievable annual revenue in the stochastic model is upper bounded up to a small error term by the optimal value of the LP, and the bound becomes relatively tight as the maximum flock size ℓ^* and the time horizon T both increase. Moreover, we construct an operating policy that provably attains annual revenue close to the optimal value of the LP under large flock sizes and long time horizons. Importantly, these results do not depend on the assumptions of Theorem 1; they hold under general price structures, birth rates, and loss rates.

These results are important because they confirm that the LP analysis provides meaningful guidance for flock management. We further buttress this conclusion with numerical experiments showing that, under realistic parameter values, the policy constructed to establish asymptotic optimality indeed approaches the LP optimum as flock sizes increase.

This policy is designed to facilitate analytical guarantees and asymptotic results, rather than to optimize performance for smaller flocks or shorter horizons. Accordingly, we also evaluate a simpler model-predictive control (MPC) policy. In this approach, at the end of each year, we solve a modified version of the LP (2) that decides on immediate sales decisions along with flock management decisions for the following years. We then implement the immediate sales decisions from the optimal solution, and then re-optimize in the following year. Numerical results demonstrate that this MPC policy performs extremely well across a broad range of practical parameter settings.

For period $t = 0, 1, 2, \ldots$, let $(\mathbf{Y}(t), \mathbf{Z}(t), \mathbf{L}(t), N(t))$ give the vector, $\mathbf{Y}(t)$, of age-specific counts of the flock (indexed by $i = 1, 2, \ldots, n$, henceforth called the *census*), the vector, $\mathbf{Z}(t)$, of sales counts (of ewe lambs and ewes indexed by $i = 0, 1, 2, \ldots, n$), the vector, $\mathbf{L}(t)$, giving the counts of lost ewes (indexed by age, $i = 0, 1, 2, \ldots, n$, where $\mathbf{L}_0(t)$ is the number of lost ewes) and a scalar, $N(t)$, giving the number of ewe lamb births of the flock in time step t. This notation reflects the fact that these quantities are random variables resulting from the application of a sales policy, as well as random lamb yields and annual losses. The policy is non-anticipating, so that $\mathbf{Z}(t)$ is determined at time t based only on the information observed to time $t-1$, along with $\mathbf{Y}(t), N(t)$ and $\mathbf{L}(t)$. For some number of years, $T \geq 1$, define $\bar{\mathbf{Y}}(T) = T^{-1} \sum_{t=0}^{T-1} \mathbf{Y}(t)$ as the average census, with similar definitions for $\bar{\mathbf{L}}(T), \bar{N}(T)$ and $\bar{\mathbf{Z}}(T)$. Also define $S(t) = \sum_{i=1}^{n} \mathbf{Y}_i(t)$ as the total count of ewes at time t and $\bar{S}(T)$ as the average total count of ewes over the first T periods.

5.1. No Policy Can Beat the Linear Program, Asymptotically

We first show that, under reasonable conditions, no policy can attain a long-run (for large T) revenue that beats the optimal objective function value of the LP (4) by an appreciable margin, at least for large flock sizes ℓ^*. The optimal value $\rho(\ell^*)$ of the LP (4) is simply $\ell^* \rho(1)$, i.e., exactly linear in ℓ^*, as follows by noting that the feasible region of (4) scales exactly linearly with ℓ^*.

We assume the following.

A1. $\mathbf{Y}(t) \geq 0, \mathbf{Z}(t) \geq 0$.

A2. The flock size constraint is met at least asymptotically, i.e., there is a function g_1 with the property that $g_1(y) = o(y)$ as $y \to \infty$ such that

$$\ell^* - g_1(\ell^*) \leq \liminf_{T\to\infty} \bar{S}(T) \leq \limsup_{T\to\infty} \bar{S}(T) \leq \ell^* + g_1(\ell^*) \quad \text{a.s.} \tag{8}$$

This constraint reflects the fact that we cannot exactly match the desired flock size at all times due to random lamb yields and random losses. Still, the constraint needs to be met in some sense, here approximately in ℓ^* asymptotically as $T \to \infty$.

A3. The number of ewes, $S(t)$, is uniformly bounded at all times t by some deterministic constant $K = K(\ell^*) \leq c_1 \ell^*$ for some $c_1 > 0$. This bound can be achieved through sales.

A4. Conditional on the number of age-i ewes, $\mathbf{Y}_i(t), \mathbf{L}_i(t)$ is binomially distributed with parameters $\mathbf{Y}_i(t)$ and the loss rate ψ_i independently of all else, for $i = 1, 2, \ldots, n$.

A5. Conditional on the number of births, $N(t)$, the number of lost lambs, $\mathbf{L}_0(t)$, is binomially distributed with parameters $N(t)$ and the loss rate ψ_0 independently of all else.

A6. Conditional on the census, $\mathbf{Y}(t)$, the number of births, $N(t)$ is distributed as a nonnegative integer-valued random variable that is bounded by $c_2 S(t) = c_2 \sum_{i=1}^{n} \mathbf{Y}_i(t)$ for some $c_2 > 0$, independent of all else. For sheep it is reasonable to take $c_2 = 4$.

A7. Age n sheep are sold after any losses, i.e., $\mathbf{Y}_n(t) - \mathbf{L}_n(t) = \mathbf{Z}_n(t)$ for all $t \geq 0$.

A8. We have conservation of flow of ewes. Most of the time we have the flow constraint that for all $t \geq 0$ and all $i = 2, 3, \ldots, n$, $\mathbf{Y}_{i-1}(t) - \mathbf{L}_{i-1}(t) - \mathbf{Z}_{i-1}(t) = \mathbf{Y}_i(t+1)$. The exception is that there is always a (very) small chance that all sheep and lambs are lost at time t, i.e., that the event

$$D(t) = \{\mathbf{L}_i(t) = \mathbf{Y}_i(t), i = 1, 2, \ldots, n \text{ and } \mathbf{L}_0(t) = N(t)\}$$

occurs. Without some way to restore the flock, it would be impossible to meet the flock size constraint, even asymptotically. Thus, we need a mechanism to restore the flock when the event $D(t)$ happens. When $D(t)$ happens, we assume that the flock is restored to the level $a = (a_1, a_2, \ldots, a_n)$ at time $t+1$, e.g., by obtaining sheep at auction, where $a_i \geq 1$ for at least one $i = 1, 2, \ldots, n$. Thus, conservation of flow can be expressed as

$$\begin{aligned}\mathbf{Y}_i(t+1) &= [\mathbf{Y}_{i-1}(t) - \mathbf{L}_{i-1}(t) - \mathbf{Z}_{i-1}(t)](1 - \mathbb{I}(D(t))) + a_i \mathbb{I}(D(t)) \\ &= \mathbf{Y}_{i-1}(t) - \mathbf{L}_{i-1}(t) - \mathbf{Z}_{i-1}(t) + a_i \mathbb{I}(D(t)) \end{aligned} \tag{9}$$

for $i = 2, 3, \ldots, n$ and all $t \geq 0$, since when $D(t)$ happens, there are no sheep to sell and so $\mathbf{Z}_{i-1}(t) = 0$. We assume that the vector a does not depend on ℓ^*.

A9. We have conservation of flow of ewe lambs, i.e., for all $t \geq 0$,

$$\mathbf{Y}_1(t+1) = N(t) - \mathbf{L}_0(t) - \mathbf{Z}_0(t) + a_1 \mathbb{I}(D(t)). \tag{10}$$

A10. The number of surviving ram lambs each year, conditional on the census in that year, has the same conditional distribution as the number of surviving ewe lambs and is conditionally independent of all else. Moreover, all ram lambs that survive the first year are sold.

In preparation for our next result, define $\|a\| = \sum_{i=1}^n |a_i|$ to be the L_1 norm. Also, recall that $\mathcal{L}^R(t)$ is the number of ram lambs that survive their first year, and let $\bar{\mathcal{L}}^R(T)$ be the average of these values over $t = 0, 1, \ldots, T-1$.

Theorem 2. *Under the assumptions A1 - A10, above, relative to the optimal objective function value $\ell^* \rho(1)$ of the steady-state LP* (4)*, there is a constant c_3 that does not depend on ℓ^* such that*

$$\limsup_{T\to\infty} \left[\sum_{i=0}^n p_i \bar{\mathbf{Z}}_i(T) + p^R \bar{\mathcal{L}}^R(T) \right] \leq \ell^* \rho(1) + c_3 (g_1(\ell^*) + \|a\|).$$

Hence,

$$\limsup_{\ell^*\to\infty} \limsup_{T\to\infty} \frac{\sum_{i=0}^n p_i \bar{\mathbf{Z}}_i(T) + p^R \bar{\mathcal{L}}^R(T)}{\ell^*} \leq \rho(1).$$

Theorem 2 shows that any policy that asymptotically maintains ℓ^* ewes cannot do much better in revenue per period than $\rho(\ell^*) = \ell^* \rho(1)$, the optimal objective function value of the LP (4). The extent to which it can exceed the optimal value of the LP is driven by fluctuations around the flock-size constraint (8) and by the need to account for occasional flock restorations due to rare events in which annual losses wipe out the entire flock. Since the restoration level $\|a\|$ can be as small as 1, this effect is very small. Effectively then, the LP provides an upper bound on the achievable performance of any policy.

5.2. An Asymptotically Optimal Policy

Thus far, we see that no policy can appreciably beat the optimal objective value of the LP, at least for large flock sizes. We now present a policy that, asymptotically as $\ell^* \to \infty$, attains the optimal objective value of the steady-state LP (4), thereby showing that the optimal value of the LP is a tight bound on achievable performance. The policy we provide here is tailored to ensure that we can *prove* that asymptotic optimality is possible. However, we do not mean to imply that one should adopt this policy in practice; the numerical results to follow indicate that a different policy performs substantially better in practical settings, though we do not have a proof of that. The asymptotically optimal policy and its proof are constructed as follows.

1. Solve the steady-state LP for a desired flock size $\ell^* = 1$ to get an optimal solution $y^* = (y_1^*, y_2^*, \ldots, y_n^*)$ that sums to 1, with optimal sales rates $z^* = (z_0^*, z_1^*, \ldots, z_n^*)$. Due to the scaling property of the feasible region, an optimal solution to the LP for any desired flock size $\ell^* > 0$ is simply $(\ell^* y^*, \ell^* z^*)$.
2. Suppose that $y_n^* > 0$, i.e., that the flock does not reduce to 0 until Year n. If not, then simply redefine n as the last year k for which $y_k^* > 0$.
3. If $z_0^* > 0$, i.e., at least some ewe lambs are sold, then define $(\tilde{y}, \tilde{z}) = (y^*, z^*)$. Otherwise, $z_0^* = 0$, in which case we modify (y^*, z^*) to obtain $(\tilde{y}, \tilde{z})$ as follows. First, recall that we defined $\zeta_0 \triangleq 1$, and for $i = 1, 2, \ldots, n$, $\zeta_i \triangleq (1-\psi_1)(1-\psi_2)\cdots(1-\psi_i)$. Define the vector $\tilde{y} = y^* + b\,(\zeta_0, \zeta_1, \zeta_2, \ldots, \zeta_{n-1})$ for some $b > 0$. This yields a total of bR_n additional ewes relative to a flock of size 1, where $R_n = \sum_{i=0}^{n-1} \zeta_i$; we will eventually choose b to be small relative to ℓ^* and positive. These additional ewes will have an expected number of non-lost ewe lambs equal to $b(1-\psi_0)\sum_{j=1}^{n} \alpha_j \zeta_{j-1}$ each year. Under the stability condition $(1-\psi_0)\sum_{j=1}^{n} \alpha_j \zeta_{j-1} > 1$, in expectation we can then sell $\tilde{z}_0 = b(1-\psi_0)\sum_{j=1}^{n} \alpha_j \zeta_{j-1} - b > 0$ non-lost ewe lambs each year, while we set the other sales rates $\tilde{z}_k = z_k^*$ for $k = 1, 2, \ldots, n-1$ and $\tilde{z}_n = (1-\psi_n)\tilde{y}_n$. Thus, we have constructed a pair $(\tilde{y}, \tilde{z})$ that is close to (y^*, z^*) with $\tilde{z}_0 > 0$.
4. Fix $\epsilon > 0$ and define the *target state* as the vector of age-specific population levels

$$\xi = (\xi_1, \xi_2, \ldots, \xi_n) = \lceil \ell^* (\tilde{y}_1, \tilde{y}_2 - \epsilon, \tilde{y}_3 - 2\epsilon, \ldots, \tilde{y}_n - (n-1)\epsilon) \rceil, \tag{11}$$

   where the ceiling operator $\lceil \cdot \rceil$ is applied elementwise. The purpose of the ϵ terms is to provide some "slack" so that the probability of returning to the state ξ starting from that state is very close to 1, due to our sales policy and concentration inequalities, as detailed next.

5. Consider the (stochastic) policy that, starting from any state, $\mathbf{Y}$, observed after the (stochastic) births of lambs and losses of lambs and ewes of all ages, sells lambs and ewes so as to bring the census (i.e., the vector giving the number of ewes at each age) for the next year down to the target state, while also selling all ewes of age n. More precisely, the policy sells all $\mathbf{Y}_n$ age-n ewes, and for $i = 0, 1, \ldots, n-1$, sells $\max\{0, \mathbf{Y}_i - \xi_{i+1}\}$ ewe lambs ($i = 0$) or ewes ($i = 1, 2, \ldots, n-1$). The difference of indices (i versus $i+1$) arises because, after sales, the time index is increased by 1. All ram lambs are sold. Recall that we assume that if ever losses completely wipe out the flock, then the flock is restored to the levels $a = (a_1, a_2, \ldots, a_n)$ with $a_k \geq 1$ for at least one k.

The stated policy induces a Markov chain $(Y(t) : t \geq 0)$ on the state space $\mathcal{S} = (y \in \mathbb{Z}^n : 0 \leq y_i \leq \xi_i, i = 1, 2, \ldots, n, y \neq 0)$. Pathwise, this chain remains in the target state for very long stretches of time, punctuated by only brief excursions out of this state. We then use a regenerative process argument to show that the policy both maintains a flock size that is approximately equal to ℓ^* and attains a near-optimal objective function value. We prove these observations below, but first we need some additional assumptions.

A11. Let $W_{ij}(t)$ be the number of non-lost (surviving for at least 1 year) ewe lambs born by the j-th ewe of age i years at time t. We assume that, for each $i = 1, 2, \ldots, n$, $W_i = (W_{ij}(t) : t \geq 0, j \geq 1)$ is a collection of i.i.d. bounded random variables with mean $\alpha_i(1 - \psi_0)$, and we assume that these collections are mutually independent. We ignore (do not use) the values $W_{ij}(t)$ when j is greater than the age-i ewe population, $\mathbf{Y}_i(t)$, at time t. Then, conditional on the census, $\mathbf{Y}(t)$, the number of ewe lambs born at time t that are not lost over their first year is $B(t) = \sum_{i=1}^{n} \sum_{j=1}^{\mathbf{Y}_i(t)} W_{ij}(t)$, which has conditional mean $(1 - \psi_0) \sum_{i=1}^{n} \alpha_i \mathbf{Y}_i(t)$ conditional on $\mathbf{Y}(t)$. We adopt a symmetric assumption for ram lambs, and assume that ram lamb and ewe lamb dynamics are independent, conditional on the census.

A12. Choose the parameters of the policy $\epsilon = \epsilon(\ell^*) = (\ell^*)^{-r_1}$ and $b = b(\ell^*) = (\ell^*)^{-r_2}$ with $0 < r_2 < r_1 < 1/2$.

A13. The stability condition $\eta := (1 - \psi_0) \sum_{j=1}^{n} \alpha_j \zeta_{j-1} > 1$ holds. Here, η is the mean number of ewe lambs that survive their first year, that are born to a single ewe over the course of its lifetime from Year 1 onwards.

A14. Let $V_{ij}(t)$ equal 1 if the jth sheep of age i is lost in Year t, and equal 0 otherwise. We assume that, for each $i = 0, 1, 2, \ldots, n$, $V_i = (V_{ij}(t) : t \geq 0, j \geq 1)$ is an i.i.d. collection of binary random variables with $V_{ij}(t) = 1$ with probability ψ_i and 0 otherwise. Moreover, we assume these collections are mutually independent and independent of the birth processes.

We are now ready to state the main result of this subsection.

THEOREM 3. *Assume A4-A14. Then, the mean steady-state flock size is equal to $\ell^* + o(\ell^*)$ and has optimal objective value $\rho\ell^* + o(\ell^*)$. Thus, as $\ell^* \to \infty$, the policy is asymptotically optimal.*

Proof Sketch: The full proof is given through a series of lemmas in the appendix; here we give only a sketch. The Markov chain describing the census as a function of time is regenerative, with cycles defined by successive visits to the target state. Concentration inequalities imply that, starting from the target state, the process returns to the target state with very high probability, so that cycles are almost always of length 1. On those rare occasions when the process does not immediately return to the target state, we use a geometric-trials argument to establish a uniform bound on the expected return time. The revenue earned while in the target state is very close to the objective value of the steady-state LP. Also, we show that the flock-size constraint is satisfied up to a vanishing error for large flock sizes. Taken together, these observations yield the result.

Theorem 3 establishes that the proposed policy is asymptotically optimal: as the target flock size ℓ^* grows, the policy attains, relatively speaking, the optimal value of the steady-state LP. When combined with Theorem 2, this result implies that the LP optimum provides a tight bound on the long-run performance that is achievable under stochastic dynamics.

5.3. A Numerical Exploration

We now know that the policy defined in the previous section is asymptotically optimal for large flock sizes and large time horizons. But how does it perform for practical flock sizes and time horizons?

Consider the realistic parameters given in Section 3.2. Solving the steady-state LP (4) for a maximum flock size of $\ell^* = 1$ yields the solution $y^* = (0.212, 0.206, 0.200, 0.194, 0.188)$ and $z^* = (0.661, 0, 0, 0, 0, 0.182)$ with optimal objective function value $\rho(1) = 335.6$. The optimal solution for a maximum flock size $\ell^* \neq 1$ is simply these values multiplied by ℓ^*. We simulated 30 independent and identically distributed replications of flock dynamics over a time horizon of 20 years, for maximum flock sizes ranging from $\ell^* = 100$ to $\ell^* = 6000$ ewes, under four different choices of the initial census:

1. The optimal LP solution $\ell^* y^*$, rounded to integers;
2. An even distribution of ewe ages $\ell^*(0.2, 0.2, 0.2, 0.2, 0.2)$;
3. All Year 1 ewes $(\ell^*, 0, 0, 0, 0)$; and
4. All Year 5 ewes $(0, 0, 0, 0, \ell^*)$.

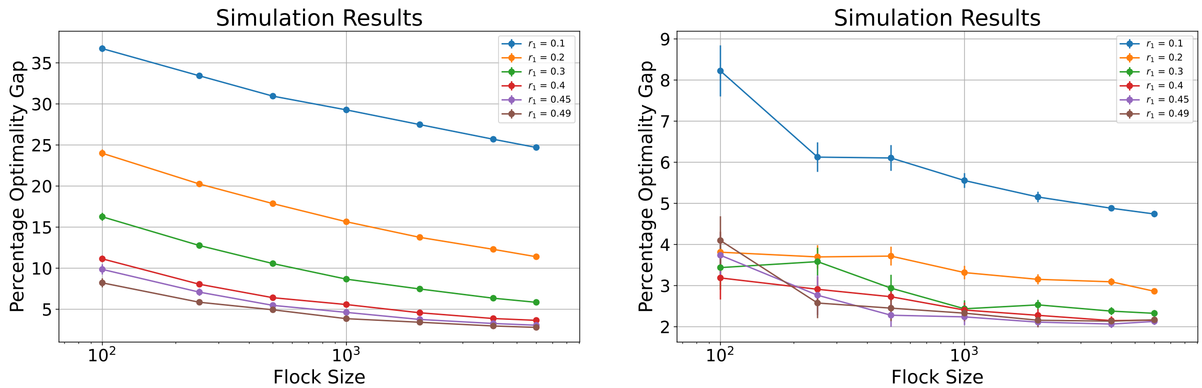

(a) Initialized from the target state (b) Renormalized to sum to ℓ^* and rounded

Figure 3 The estimated expected sub-optimality gap of various policies, indexed by r_1. The horizontal axis denotes the maximum flock size ℓ^* and is in log scale. The vertical axis gives the sub-optimality percentage of the policy relative to the objective function of the LP, measured as $100(\rho(\ell^*) - V)/\rho(\ell^*)$, where V is the sample average over replications of the simulated profits. Lower values are more desirable. Error bars denote 95% confidence intervals around each estimated value, but are difficult to see in Figure 3a because they are very small.

All four of these choices for the initial census attain the maximum flock size ℓ^*, in accord with the results in Section 3.2 that reaching the maximum flock size immediately through initial purchases is often the optimal approach. Ewes of age i ($i = 1, 2, \ldots, n$) are assumed to have a binomially distributed number of lambs with parameters 3 and $2\alpha_i/3$, so that the average number of lambs per age-i ewe is $2\alpha_i$. All lambs are independently equally likely to be ewe lambs or ram lambs. Losses of ewes and lambs through the year are assumed to be binomially distributed within each age cohort, again independently of all else.

We simulated policies as defined in Section 5.2, which depend on the two parameters r_1 and r_2. Recall that the parameter r_2 is used to construct the target state if the optimal solution (y^*, z^*) does not prescribe the sale of ewe lambs. The optimal solution, above, *does* prescribe selling ewe lambs, as does the optimal solution for all plausible input parameters that we tried. Therefore, we do not experiment with the parameter r_2 here. We simulated policies for several choices of the parameter r_1 in the interval $(0, 0.5)$, as required by the theory. This parameter is used to set the additional parameter $\epsilon(\ell^*) = (\ell^*)^{-r_1}$ used to construct the target state ξ in (11). Initial experiments found that this formula for ϵ yielded very high sales levels. Accordingly, we instead chose $\epsilon(\ell^*) = (\ell^*)^{-r_1}/20$. This simple modification does not change the conclusions of the theoretical results in Section 5.2. Thus, we adopt this choice here. The policies from Section 5.2 that we simulated involve computing the target state (11) and, at the end of each year of operations, selling ewes as described in Item 5 of that section. In our simulations, we never saw a year in which the entire flock needed to be restored to the levels given by the vector a (see Assumption A8 in Section 5.1), so we do not specify the value of that vector.

Figure 3a depicts the results assuming the initial flock census is the target state, which is the first of the four options for the initial flock census listed above. The policies perform poorly at small flock sizes; however, performance improves steadily as flock size increases, and the gap relative to the steady-state LP objective narrows, consistent with the theoretical predictions. The results indicate that the parameter r_1 should be selected close to its upper bound of 0.5, which yields smaller ϵ-adjustments to the target state (see Item 4 in Section 5.2).

These results suggest that the suboptimality of the policies is partially due to the ϵ terms that reduce the flock size below its target value. Those adjustments are relatively small asymptotically as the flock size grows, but are relatively large for small flock sizes. To test that observation, the policies in Figure 3b renormalize the target state to match the flock size ℓ^* after the ϵ terms are applied. The results improve substantially; note the change in the vertical axis scale.

We omit results from simulations that start from other flock census options (Options 2 through 4). Those results were similar, with worse performance when starting with all Year 5 ewes (Option 4) and better performance for an even age distribution (Option 2) and all Year 1 ewes (Option 3).

Taken together, these results suggest that performance improves as r_1 increases, i.e., as the target state ξ gets closer to $\lceil \ell^* y^* \rceil$ (with the ceiling operation applied element-wise). However, the variation in performance when the initial census varies over the four options is undesirable.

This fragility to the initial census motivates an additional class of policies. We used *model predictive control* (MPC) to choose the sales decisions. More precisely, whenever a sales decision is required, we use the time-dependent LP defined by (2) with no discounting ($r = 0$) and no initial purchases ($\mathbf{x} = 0$). Moreover, we add additional decision variables for the immediate sales decisions, $w = (w_0, w_1, \ldots, w_n)$, which, combined with the current census $\mathbf{Y}$, determine the vector $\mathbf{y}(0)$ that is the initial flock size in the time-dependent LP used in MPC. We do not apply losses to the census $\mathbf{Y}$ before making sales decisions, because those losses have already occurred at the time (the end of the year) when sales decisions are made. We use the notation $\tilde{T}$ to distinguish the time horizon in the MPC formulation from the time horizon $T = 20$ years over which we simulated.

We chose $\tilde{T}=5$ (results for other values of $\tilde{T}$ were similar). The full formulation of this MPC optimization problem is as follows:

$$
\begin{array}{lll}
\max & \sum_{i=0}^{n} p_i w_i + \sum_{t=0}^{\tilde{T}-1}\left[\sum_{i=0}^{n} p_i \mathbf{z}_i(t) + p^R(1-\psi_0)\sum_{i=1}^{n}\alpha_i \mathbf{y}_i(t) - c^H \sum_{i=1}^{n}\mathbf{y}_i(t)\right] & \\
\text{s/t} & (1-\psi_0)\sum_{j=1}^{n}\alpha_j \mathbf{y}_j(t-1) - \mathbf{z}_0(t-1) - \mathbf{y}_1(t) = 0 & t=1,2,\dots,\tilde{T} \\
& (1-\psi_{i-1})\mathbf{y}_{i-1}(t-1) - \mathbf{z}_{i-1}(t-1) - \mathbf{y}_i(t) = 0 & t=1,2,\dots,\tilde{T}; i=2,3,\dots,n \\
& (1-\psi_n)\mathbf{y}_n(t) - \mathbf{z}_n(t) = 0 & t=0,1,\dots,\tilde{T}-1 \\
& \sum_{i=1}^{n}\mathbf{y}_i(t) \le \ell^* & t=0,1,\dots,\tilde{T}-1 \\
& \mathbf{y}_i(0) + w_{i-1} = \mathbf{Y}_{i-1} & i=1,2,\dots,n \\
& w_n = \mathbf{Y}_n & \\
& w_i \ge 0 & i=0,1,\dots,n \\
& \mathbf{y}_i(t) \ge 0 & t=0,1,\dots,\tilde{T}; i=1,2,\dots,n \\
& \mathbf{z}_i(t) \ge 0 & t=0,1,\dots,\tilde{T}-1; i=0,1,\dots,n.
\end{array}
$$

The results appear in Figure 4, where we vary both the flock size ℓ^* and the choice of initial census. The MPC policy is within a few percentage points of the optimal LP value for all flock sizes and all initial census options. Given that we proved that one cannot do much better than the LP value, this policy has very impressive performance, even though we have not proved that the policy is asymptotically optimal. MPC even does better than the optimal LP value in some cases, e.g., when we start with a young census, which cannot be maintained in a steady-state situation. We believe that the residual sensitivity to the initial census seen in these results is due to the chosen simulation time horizon of $T=20$ years. Indeed, in results not shown here, the errors relative to the optimal steady-state LP objective were further reduced when we chose $T=40$ years. We elected not to present those results because performance over such long time scales is unlikely to be the primary driver of decisions in practice. It is not clear why the performance of MPC is worst when the initial census is the optimal solution of the LP, but perhaps that relates to the difference in perspective of the short planning horizon of MPC relative to the long-term perspective of the steady-state LP. The decrease in variability in the simulation runs as the flock size grows is evident, reflecting the asymptotic regime of our analytical results.

6. Conclusion

To help facilitate an innovative solution to a critical maintenance challenge for solar sites with photovoltaic panels, we provided an operations management perspective for how farmers can

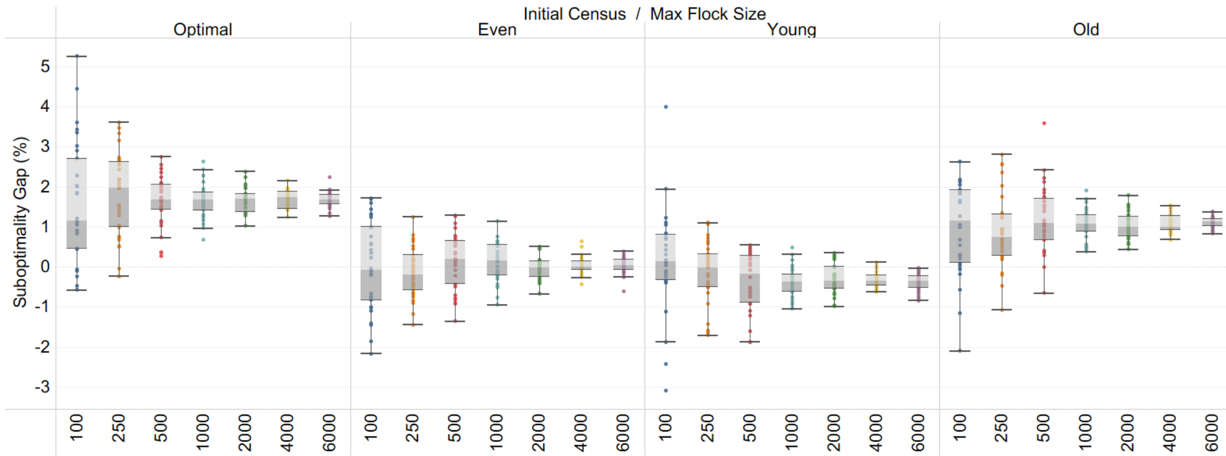

Figure 4 **Simulated sub-optimality gap of the MPC policy by flock size ℓ^* and the initial census (Options 1 through 4) expressed as a percentage. The sub-optimality percentage of the policy relative to the objective function of the LP is measured as $100(\rho(\ell^*) - V)/\rho(\ell^*)$, where V is the sample average over replications of the simulated profits. Thus, negative values indicate better performance than the steady-state LP objective. Whiskers on the graph extend to 1.5 times the interquartile range.**

optimally manage sheep flocks to control vegetation while also benefiting from lamb sales. We began by analyzing the flock growth phase, where our model demonstrates that total profit is a piecewise linear, concave function of the number of sheep initially purchased externally. We then turned to the steady-state maintenance problem, and proved that a simple yet widely used "maximum-lifetime" policy is optimal under realistic assumptions. Furthermore, we showed that randomness in birth and loss rates does not undermine this core insight in the long run: for large flocks, no policy can outperform the long-run revenue of the deterministic steady-state optimum. Finally, we constructed a class of policies that provably asymptotically achieves the deterministic model's optimal performance as the flock size grows, and provided a heuristic policy based on model-predictive control which, in numerical results, achieves near optimal performance even at smaller flock sizes.

Beyond these theoretical contributions, we partnered with a solar grazing platform to implement our findings in a user-friendly spreadsheet tool. This tool has guided dozens of farmers and helped some farmers to secure loans for initial flock procurement. The analysis has also been used in municipal approval processes for new solar sites, where town leaders can see that grazing sheep preserves — or even enhances — the agricultural value of the land under consideration.

Looking ahead, there are many opportunities for further research on solar grazing and other forms of agrivoltaics. Incorporating more nuanced cost structures, the opportunity to utilize multiple livestock species like pigs and poultry, site-specific vegetation decisions, or seasonal fluctuations in lamb prices, could lead to richer models and insights. As we have demonstrated in this work, the deeper levels of understanding that come from new operations models can unlock economic value for these more sustainable forms of food and energy production.

Acknowledgments

We express our gratitude to Caleb Scott of United Agrivoltaics for providing deep domain knowledge and for iteratively engaging with us as he made use of this research. Rithvik Duggireddy and Mingfei Liao, undergraduates at Cornell at the time, provided early modeling efforts related to the topic of this paper. Will Ma helped with early scoping and exploration. Finally, we thank a reviewer who suggested the current proof of Theorem 1; the original proof relied on a lengthy exchange argument. This work was partially supported by National Science Foundation grant OAC 2410950.

References

Ata B, Lee D, Sönmez E (2019) Dynamic volunteer staffing in multicrop gleaning operations. *Operations Research* 67(2):295–314.

Beckwith S (2021) *The Solar Grazing Planner: An Easy-to-Use Extension Tool for Sheep Producers Rotationally Grazing on Solar Sites*. Master's thesis, Cornell University.

Bertsimas D, Tsitsiklis JN (1997) *Introduction to Linear Optimization* (Belmont, MA: Athena Scientific).

Bowen MK (2024) Agrivoltaic grazing systems in Queensland: The potential, challenges, and opportunities for the coexistence of grazing and solar power plants in Queensland. Project report, Department of Agriculture and Fisheries, State of Queensland, URL `https://era.dpi.qld.gov.au/id/eprint/12988/`, accessed April 24, 2025.

Bu J, Gong X, Chao X (2023) Asymptotic optimality of base-stock policies for perishable inventory systems. *Management Science* 69(2):846–864.

Chao X, Gong X, Shi C, Zhang H (2015) Approximation algorithms for perishable inventory systems. *Operations Research* 63(3):585–601.

Chen H, Ryan JK (2023) Optimal specialty crop planning policies with yield learning and forward contract. *Production and Operations Management* 32(2):359–378.

Chen X, Pang Z, Pan L (2014) Coordinating inventory control and pricing strategies for perishable products. *Operations Research* 62(2):284–300.

Dahl G, Hjort A (1976) *Having Herds: Pastoral Herd Growth and Household Economy*, volume 2 of *Stockholm Studies in Social Anthropology* (Stockholm: Department of Social Anthropology, University of Stockholm).

FarmingAhead (2024) Historic price disparity emerges between mutton and lamb. URL `https://www.farmingahead.com.au/livestock/news/1463441/historic-price-disparity-emerges-mutton-lamb`, accessed: 2024-09-21.

Farrell L, Kenyon P, Tozer P, Ramilan T, Cranston L (2020) Quantifying sheep enterprise profitability with varying flock replacement rates, lambing rates, and breeding strategies in New Zealand. *Agricultural Systems* 184:102888.

Feldman D, Zuboy J, Dummit K, Stright D, Heine M, Grossman S, Margolis R (2024) Spring 2024 solar industry update. Technical report, National Renewable Energy Laboratory (NREL), Golden, CO (United States).

Flay KJ, Ridler AL, Compton CW, Kenyon PR (2021) Ewe wastage in New Zealand commercial flocks: extent, timing, association with hogget reproductive outcomes and BCS. *Animals* 11(3):779.

Fonsêca VdFC, Culhari EdA, Moura GAB, Nascimento ST, Milan HM, Chiquitelli Neto M, Maia ASC (2023) Shade of solar panels relieves heat load of sheep. *Applied Animal Behaviour Science* 265:105998.

Gasch A, Lara R, Pearce JM (2025) Financial analysis of agrivoltaic sheep: Breeding and auction lamb business models. *Applied Energy* 381:125057.

Hoffman AJ (1952) On approximate solutions of systems of linear inequalities. *Journal of Research of the National Bureau of Standards* 49(4):263–265.

Hu B, Sun Z (2022) Managing self-replicating innovative goods. *Management Science* 68(1):399–419.

Karaesmen IZ, Scheller-Wolf A, Deniz B (2011) Managing perishable and aging inventories: review and future research directions. *Planning Production and Inventories in the Extended Enterprise: A State of the Art Handbook, Volume 1* 393–436.

Kochendoerfer N, Hain A, Thonney M (2019) The agricultural, economic and environmental potential of co-locating utility scale solar with grazing sheep. *Atkinson Center for a Sustainable Future, Cornell University Ithaca, NY* .

Kochendoerfer N, Thonney ML (2021) Grazing sheep on solar sites in New York State: Opportunities and challenges. *Scope and scaling-up of the NYS sheep industry to graze ground-mounted photovoltaic arrays for vegetation management., Cornell University Atkinson Center for a Sustainable Future, Ithaca, NY* .

Kouvelis P, Liu Y, Qiu Y, Turcic D (2023) Managing operations of a hog farm facing volatile markets: Inventory and selling strategies. *Manufacturing & Service Operations Management* 25(5):1711–1729.

Levi R, Rajan M, Singhvi S, Zheng Y (2024) Improving farmers' income on online agri-platforms: Evidence from the field. *Manufacturing & Service Operations Management* 26(2):345–360.

Mazinani M, Rude B (2020) Population, world production and quality of sheep and goat products. *American Journal of Animal and Veterinary Sciences* 15(4):291–299.

Nahmias S (2011) *Perishable inventory systems*, volume 160 (Springer Science & Business Media).

Notter D (2000) Effects of ewe age and season of lambing on prolificacy in US Targhee, Suffolk, and Polypay sheep. *Small Ruminant Research* 38(1):1–7.

Ontario Sheep Marketing Association (2012) Ontario sheep economic workbook: Annual spring lambing flock. `https://www.ontariosheep.org/media/gjzduads/osma-sheep-economics-workbook-annual-lambing-1.pdf`, accessed: 2024-08-06.

OSU Sheep Team (2018) Cull ewe checklist. URL `https://u.osu.edu/sheep/2018/12/18/cull-ewe-checklist/`, [Accessed: 2024-09-22].

Peng X, Wu OQ, Souza GC (2024) Renewable, flexible, and storage capacities: Friends or foes? *Manufacturing & Service Operations Management* 26(5):1730–1749.

Plambeck EL (2013) OM Forum—Operations management challenges for some “cleantech” firms. *Manufacturing & Service Operations Management* 15(4):527–536.

Resnick SI (2013) *Adventures in Stochastic Processes* (Springer Science & Business Media).

Ross SM (1996) *Stochastic Processes* (New York: Wiley), 2nd edition.

Secomandi N (2025) OM Forum—Asset-level perspectives on the clean energy transition. *Manufacturing & Service Operations Management* 27(1):8–20.

Shiflett JS (2022) Mount Morris agrivoltaic study. Technical report, Agrivoltaic Solutions, URL `https://solargrazing.org/wp-content/uploads/2022/01/MountMorris-AgrivoltaicReport-FINAL_PRINT_ready.pdf`.

Sorensen A, Nogeire T, Hunter M, Beck S, Hunter M, Murphy R, Sorensen A (2022) Potential placement of utility-scale solar installations on agricultural lands in the US to 2040. Technical report, American Farmland Trust.

Stygar A, Kristensen AR, Makulska J (2014) Optimal management of replacement heifers in a beef herd: A model for simultaneous optimization of rearing and breeding decisions. *Journal of Animal Science* 92(8):3636–3649.

Talluri KT, Van Ryzin GJ (2006) *The theory and practice of revenue management*, volume 68 (Springer Science & Business Media).

Thorne J, Murdoch B, Freking B, Redden R, Murphy TW, Taylor J, Blackburn HD (2021) Evolution of the sheep industry and genetic research in the United States: opportunities for convergence in the twenty-first century. *Animal Genetics* 52(4):395–408.

US Department of Agriculture Agricultural Marketing Service (2025) Sheep auction reports. URL `https://www.ams.usda.gov/market-news/sheep-auction-reports`, accessed: 2025-03-29.

USDA Economic Research Service (2025) Sheep, lamb & mutton - sector at a glance. URL `https://www.ers.usda.gov/topics/animal-products/sheep-lamb-mutton/sector-at-a-glance`, accessed: 2025-03-29.

Velarde Morales JI, Xin L (2026) Agricultural index insurance design: An optimization approach. Working Paper.

Wang C, Wu Q, Pan K, Shen ZJM (2024) Layout optimization for a large-scale grid-connected solar power plant. *INFORMS Journal on Computing* .

Yan M, Schmit TM, Baker MJ, LeRoux MN, Gómez MI (2022) Sell now or later? a decision-making model for feeder cattle selling. *Agricultural and Resource Economics Review* 51(2):343–360.

Yi Z, Wang Y, Chen YJ (2021) Financing an agricultural supply chain with a capital-constrained smallholder farmer in developing economies. *Production and Operations Management* 30(7):2102–2121.

Zhang W, Gao L, Zolghadr M, Jian D, ElHafsi M (2023) Dynamic incentives for sustainable contract farming. *Production and Operations Management* 32(7):2049–2067.

Electronic Companion

EC.1. State Transition Diagram for Section 4

Figure EC.1 displays a state transition diagram for the age-based model corresponding to Section 4.

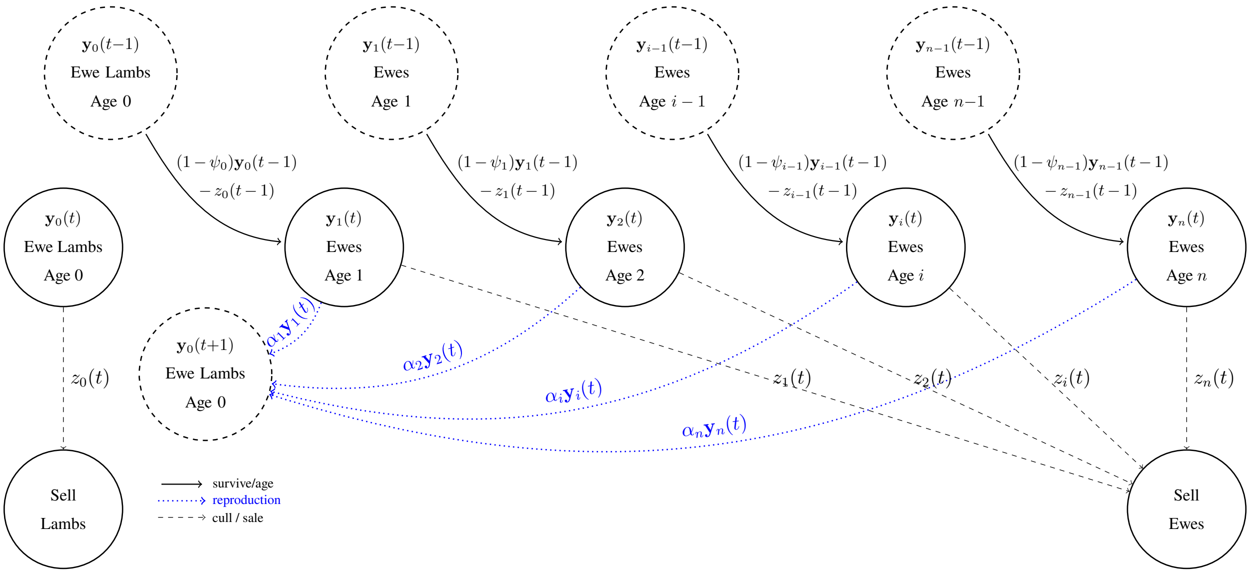

Figure EC.1 **Annual state–transition diagram of ewes in the flock for the deterministic age-structured ewe model (uniform node size). For convenience, let $\mathbf{y}_0(t+1) := \sum_{j=1}^{n} \alpha_j \mathbf{y}_j(t)$.**

EC.2. Proofs from Sections 3 and 4

Proof of Proposition 1. The function $g(\mathbf{x})$ is the sum of the linear, and thus concave, term $-\sum_{i=1}^{n} c_i \mathbf{x}_i$ and the optimal objective function value of the LP (2) where we omit the $\mathbf{x}$ term from the objective function. The condition in that LP that $\mathbf{x} \geq 0$ is simply a restriction of the domain of $\mathbf{x}$. Thus, we can view the LP (2) without the linear term in $\mathbf{x}$ in the objective function as a standard LP

$$
\begin{aligned}
\max \quad & c^T u \\
\text{s/t} \quad & Au \leq b \\
& u \geq 0,
\end{aligned}
$$

where the initial procurement vector $\mathbf{x}$ enters only through the right-hand-side vector b. By standard results in LP, the optimal value is a concave and piecewise affine function of the right-hand-side vector b, e.g., see Chapter 5.2 of Bertsimas and Tsitsiklis 1997. ■

Proof of Proposition 2. Let $(\mathbf{y},\mathbf{z})$ be a feasible solution for the full formulation (3). Define

$$(\bar{y},\bar{z}) = \left(T^{-1}\sum_{t=0}^{T-1}\mathbf{y}(t), T^{-1}\sum_{t=0}^{T-1}\mathbf{z}(t)\right)$$

as the average of the time-dependent variables over the full cycle. Then, $(\bar{y},\bar{z})$ is feasible for the steady-state formulation (4). To see why this holds, consider just the first constraint in (2). Averaging over $t = 1,2,\ldots,T$ yields

$$(1-\psi_0)\sum_{j=1}^{n}\alpha_j\bar{y}_j - \bar{z}_0 = \frac{1}{T}\sum_{t=1}^{T}\mathbf{y}_1(t) = \frac{1}{T}\sum_{t=0}^{T-1}\mathbf{y}_1(t) = \bar{y}_1,$$

because of the constraint that $\mathbf{y}_1(T) = \mathbf{y}_1(0)$. A similar argument holds for other constraints. Moreover, the objective function values in the two formulations match. Likewise, given a feasible solution for (4), (y,z), simply define $\mathbf{y}(t) = y$ for $0 \le t \le T$ and $\mathbf{z}(t) = z$ for $0 \le t < T$, which is then a feasible solution for (2) with the same objective function value. Thus, the optimal objective function value in the two formulations match. ■

Proof of Theorem 1. In the LP formulation (4), the decisions $\{z_j\}_{j=0}^{n}$ are fully determined by $\{y_j\}_{j=1}^{n}$, because both the objective function and constraints can be rewritten solely in terms of y, yielding the following reformulation:

$$\max_{y_1,\cdots,y_n} \sum_{i=1}^{n}\Big[\left(p_0+p^R\right)(1-\psi_0)\,\alpha_i + p_i\,(1-\psi_i) - p_{i-1}\Big]y_i$$

$$\text{s/t} \quad (1-\psi_0)\sum_{j=1}^{n}\alpha_j y_j \ge y_1, \tag{EC.1}$$

$$(1-\psi_{i-1})y_{i-1} \ge y_i, \quad i = 2,3,\ldots,n, \tag{EC.2}$$

$$\sum_{i=1}^{n} y_i = \ell^*, \tag{EC.3}$$

$$y_i \ge 0, \quad i = 1,2,\ldots,n.$$

Let us ignore constraint (EC.1) for the moment. Under the assumptions that $\{p_i\}_{i=0}^{n}$ are decreasing and convex, $\{\alpha_i\}_{i=1}^{n}$ are increasing, and $\{\psi_i\}_{i=0}^{n}$ are decreasing, one can verify that the coefficients $\left\{\left(p_0+p^R\right)(1-\psi_0)\,\alpha_i + p_i\,(1-\psi_i) - p_{i-1}\right\}_{i=1}^{n}$ in the objective function are increasing in i. Thus, all constraints (EC.2) are tight at the optimal solution y^*, because it is optimal to allocate as much as possible to y_i with a larger coefficient, and $\{y_i^*\}_{i=1}^{n}$ can be determined from (EC.3). Finally, under the assumption that $(1- \psi_0)\sum_{j=1}^{n}\alpha_j\zeta_{j-1} - 1 \ge 0$, the previously obtained solution $\{y_i^*\}_{i=1}^{n}$ also satisfies (EC.1), making it optimal for the full problem and yielding $z_1^* = \ldots = z_{n-1}^* = 0$.

Since $z_1^* = \cdots = z_{n-1}^* = 0$, we have $y_i^* = (1-\psi_{i-1})y_{i-1}^*$ for each $i = 2, \ldots, n$. Recalling that $\zeta_0 = 1$ and for $i = 1, 2, \ldots, n$, $\zeta_i = (1-\psi_1)(1-\psi_2)\cdots(1-\psi_i)$ it follows that

$$\begin{aligned}\ell^* &= y_1^* + y_2^* + \ldots + y_n^* \\ &= y_1^* \left[\zeta_0 + \zeta_1 + \cdots + \zeta_{n-1}\right] \\ &= y_1^* R_n,\end{aligned}$$

which implies $y_1^* = R_n^{-1}\ell^*$. Thus,

$$\begin{aligned}z_0^* &= (1-\psi_0)\sum_{j=1}^n \alpha_j y_j^* - y_1^* = \left((1-\psi_0)\sum_{j=1}^n \alpha_j \zeta_{j-1} - 1\right) R_n^{-1}\ell^*, \\ z_n^* &= (1-\psi_n)y_n^* = \zeta_n y_1^* = \zeta_n R_n^{-1}\ell^*,\end{aligned}$$

completing the proof. ■

Proof of Proposition 3. We first bound the optimal LP value from above. From (4),

$$\begin{aligned}\sum_{j=0}^n p_j z_j + p^R(z_0 + y_1) &\le p_0 \sum_{j=0}^n z_j + p_0(1-\psi_0)\sum_{j=1}^n \alpha_j y_j \\ &= p_0\left[\left(2(1-\psi_0)\sum_{j=1}^n \alpha_j y_j - y_1\right) + ((1-\psi_1)y_1 - y_2) + \cdots \right. \\ &\qquad \left. + ((1-\psi_{n-1})y_{n-1} - y_n) + (1-\psi_n)y_n\right] \\ &= p_0\left[2(1-\psi_0)\sum_{j=1}^n \alpha_j y_j - \sum_{j=1}^n \psi_j y_j\right] \\ &\le p_0\left[\max_{j=1,\ldots,n}\{2(1-\psi_0)\alpha_j - \psi_j\}\sum_{j=1}^n y_j\right] \\ &= p_0 \max_{j=1,\ldots,n}\{2(1-\psi_0)\alpha_j - \psi_j\}\,\ell^*.\end{aligned}$$

In addition, the value of the ML policy equals

$$\begin{aligned}p_0 z_0 + p_n z_n + p^R(z_0 + y_1) &\ge p_0\left(2z_0 + y_1\right) \\ &= p_0\left(2(1-\psi_0)\sum_{j=1}^n \alpha_j \zeta_{j-1} - 1\right) R_n^{-1}\ell^*,\end{aligned}$$

where we use the fact that Theorem 1 gives an expression for z_0 and its proof establishes that $y_1 = R_n^{-1}\ell^*$. The two bounds, above, yield the result. ■

EC.3. Proofs from Section 5

Proof of Theorem 2. For $t \geq 0$, let $\mathcal{F}_t$ be the sigma field generated by $(\mathbf{Y}(k), \mathbf{L}(k), N(k), \mathbf{Z}(k), \mathcal{L}^R(k) : k = 0, 1, \ldots, t)$ and let $\mathcal{F} = (\mathcal{F}_t : t \geq 0)$ be the associated filtration. Fix any $i \in \{1, 2, \ldots, n\}$. For $t \geq 0$, define $M_t = \sum_{k=0}^{t} (\psi_i \mathbf{Y}_i(k) - \mathbf{L}_i(k))$. Then, $M = (M_t : t \geq 0)$ is a martingale with respect to the filtration $\mathcal{F}$. Our assumptions ensure that the martingale differences are bounded in absolute value by K, and hence we can apply Azuma's inequality, e.g., Ross (1996, p. 307), to get that for any $T \geq 1$ and any $\delta > 0$,

$$\begin{aligned}\mathbb{P}\left(|\psi_i \bar{\mathbf{Y}}_i(T) - \bar{\mathbf{L}}_i(T)| > \delta\right) &= \mathbb{P}(|M_{T-1}| > \delta T) \\ &\leq 2 \exp\left(\frac{-\delta^2 T}{2K^2}\right).\end{aligned}$$

The right-hand side of this bound is summable in T, and thus by the Borel-Cantelli Lemma (see, e.g., Ross 1996, p. 4),

$$\psi_i \bar{\mathbf{Y}}_i(T) - \bar{\mathbf{L}}_i(T) \to 0 \tag{EC.4}$$

as $T \to \infty$ a.s. Combining this fact for $i = n$ with our assumption that age-n sheep are sold, which implies that $\bar{\mathbf{Y}}_n(T) - \bar{\mathbf{L}}_n(T) = \bar{\mathbf{Z}}_n(T)$, gives

$$(1 - \psi_n) \bar{\mathbf{Y}}_n(T) - \bar{\mathbf{Z}}_n(T) \to 0 \tag{EC.5}$$

as $T \to \infty$ a.s. In other words, the third constraint of the LP (4) holds asymptotically.

Now, by the conservation of the flow of ewes from (9), for $i = 2, 3, \ldots, n$, we have

$$\bar{\mathbf{Y}}_{i-1}(T) - \bar{\mathbf{L}}_{i-1}(T) - \bar{\mathbf{Z}}_{i-1}(T) + a_i \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{I}(D(t)) - \bar{\mathbf{Y}}_i(T) + \frac{\mathbf{Y}_i(0) - \mathbf{Y}_i(T)}{T} = 0.$$

We therefore get

$$\limsup_{T \to \infty} \left| \bar{\mathbf{Y}}_{i-1}(T) - \bar{\mathbf{L}}_{i-1}(T) - \bar{\mathbf{Z}}_{i-1}(T) - \bar{\mathbf{Y}}_i(T) \right| \in [0, a_i].$$

Then, using (EC.4), for $i = 2, 3, \ldots, n$, almost surely

$$\limsup_{T \to \infty} \left| (1 - \psi_{i-1}) \bar{\mathbf{Y}}_{i-1}(T) - \bar{\mathbf{Z}}_{i-1}(T) - \bar{\mathbf{Y}}_i(T) \right| \in [0, a_i].$$

We can show, using the same martingale argument above, that

$$\sum_{j=1}^{n} \alpha_j \bar{\mathbf{Y}}_j(T) - \bar{N}(T) \to 0, \text{ and}$$

$$\psi_0 \sum_{j=1}^{n} \alpha_j \bar{\mathbf{Y}}_j(T) - \bar{\mathbf{L}}_0(T) \to 0$$

as $T \to \infty$ a.s. Combining these observations with conservation of ewe lambs in (10) gives

$$\limsup_{T\to\infty} \left| (1-\psi_0) \sum_{j=1}^{n} \alpha_j \bar{\mathbf{Y}}_j(T) - \bar{\mathbf{Z}}_0(T) - \bar{\mathbf{Y}}_1(T) \right| \in [0, a_1]$$

as $T \to \infty$ a.s. The same argument for ewe lambs can be applied to ram lambs to establish that

$$\limsup_{T\to\infty} \left| (1-\psi_0) \sum_{j=1}^{n} \alpha_j \bar{\mathbf{Y}}_j(T) - \bar{\mathcal{L}}^R(T) \right| \in [0, a_1],$$

which implies that

$$\limsup_{T\to\infty} \left| p^R(\bar{\mathbf{Z}}_0(T) + \bar{\mathbf{Y}}_1(T)) - p^R \bar{\mathcal{L}}^R(T) \right| \le p^R a_1. \tag{EC.6}$$

We have thus shown that for $i = 2, 3, \ldots, n$, almost surely,

$$\limsup_{T\to\infty} \left| (1-\psi_{i-1}) \bar{\mathbf{Y}}_{i-1}(T) - \bar{\mathbf{Z}}_{i-1}(T) - \bar{\mathbf{Y}}_i(T) \right| \in [0, a_i] \quad \text{and} \tag{EC.7}$$

$$\limsup_{T\to\infty} \left| (1-\psi_0) \sum_{j=1}^{n} \alpha_j \bar{\mathbf{Y}}_j(T) - \bar{\mathbf{Z}}_0(T) - \bar{\mathbf{Y}}_1(T) \right| \in [0, a_1]. \tag{EC.8}$$

From (EC.6), the contribution to revenue from ram lamb sales coincides with the corresponding term in (4) almost surely as $T \to \infty$. From (8), (EC.5), (EC.7) and (EC.8), almost surely as $T \to \infty$, $(\bar{\mathbf{Y}}(T), \bar{\mathbf{Z}}(T))$ satisfies the constraints of the LP (4) except for an error in the right-hand side values that is bounded by $g_1(\ell^*) + \|a\|$. Hoffman's inequality (Hoffman 1952) then implies that, almost surely as $T \to \infty$, $(\bar{\mathbf{Y}}(T), \bar{\mathbf{Z}}(T))$ is at most a distance $c_0(g_1(\ell^*) + \|a\|)$ away from the feasible region of the LP for some constant c_0 that does not depend on ℓ^*. Thus, since the objective function is linear, we conclude that, almost surely as $T \to \infty$, the objective function of $(\bar{\mathbf{Y}}(T), \bar{\mathbf{Z}}(T))$ is at most the optimal objective value of the LP plus $c_3(g_1(\ell^*) + \|a\|)$, where

$$c_3 = c_0 \max \left\{ \max_{i=0}^{n} p_i,\ p^R \right\}.$$

■

The proof of Theorem 3 requires a series of lemmas. Our first result proves that, starting from the target state, there is a very large probability that we will return to the target state in one transition.

LEMMA EC.1. *Assume A4-A14. Then, there exist $\theta_1, \theta_2 > 0$ and a constant $\theta_3 \in (0,1)$ that do not depend on ℓ^* such that*

$$\mathbb{P}(Y(1) = \xi | Y(0) = \xi) \geq 1 - \theta_1 \exp(-\theta_2 (\ell^*)^{\theta_3})$$

for sufficiently large ℓ^*.

Proof: Fix $\theta_3 \in (0, 1-2r_1)$, where r_1 was chosen in Assumption A11. Suppose $Y(0) = \xi$. For each $i = 1, 2, \ldots, n-1$, the number of age-i ewes that are lost, L_i say, has a binomial distribution with parameters $\xi_i = \lceil \ell^*(\tilde{y}_i - (i-1)\epsilon) \rceil$ and ψ_i, and thus mean $\psi_i \lceil \ell^*(\tilde{y}_i - (i-1)\epsilon) \rceil$. Suppose first that $z_0^* = 0$ so that $\tilde{y}_i = y_i^* + b\zeta_{i-1}$. Using the facts that y^* is feasible for (4) and the fact that for any real numbers x_1 and x_2, $(1-\psi_i)(\lceil x_1 \rceil - x_1) - (\lceil x_2 \rceil - x_2) \geq -1$,

$$\begin{aligned}
&\mathbb{P}(Y_{i+1}(1) = \xi_{i+1} | Y(0) = \xi) \\
&= \mathbb{P}(L_i \leq \xi_i - \xi_{i+1} | Y_i(0) = \xi_i) \\
&= \mathbb{P}(L_i - \psi_i \xi_i \leq (1-\psi_i)\xi_i - \xi_{i+1} | Y_i(0) = \xi_i) \\
&= \mathbb{P}\left(L_i - \psi_i \xi_i \leq (1-\psi_i)\lceil \ell^*(y_i^* + b\zeta_{i-1} - (i-1)\epsilon) \rceil - \lceil \ell^*(y_{i+1}^* + b\zeta_i - i\epsilon) \rceil | Y_i(0) = \xi_i\right) \\
&\geq \mathbb{P}(L_i - \psi_i \xi_i \leq \ell^* z_i^* + \ell^* \epsilon [i - (1-\psi_i)(i-1)] - 1 | Y_i(0) = \xi_i) \\
&\geq \mathbb{P}(L_i - \psi_i \xi_i \leq \ell^* \epsilon - 1 | Y_i(0) = \xi_i) \\
&\geq 1 - \exp\left(\frac{-2(\ell^* \epsilon - 1)^2}{\lceil \ell^*(\tilde{y}_i - (i-1)\epsilon) \rceil}\right) \\
&\geq 1 - \exp(-\beta_i (\ell^*)^{\theta_3})
\end{aligned}$$

for an appropriately defined constant β_i, where the second-to-last step uses Hoeffding's inequality. The same argument applies for the case when $z_0^* > 0$, except that the terms involving b that cancel, above, are not present.

Next, starting from the target state, the number of Year 1 ewes, $Y_1(1)$, in one transition, will equal its target value provided that B, the number of ewe lambs that are born and survive their first year, is sufficiently large. The random variable B is a sum over all ewes of the (bounded by c_2) random number of their ewe offspring that survive their first year, and has mean $\mu_B := (1-\psi_0)\sum_{i=1}^n \alpha_i \xi_i$. In what follows, we again apply Hoeffding's inequality to show that B is sufficiently large with high probability. For the case where $z_0^* = 0$, using the facts that $\xi_i = \lceil \ell^*(y_i^* + b\zeta_{i-1} - (i-1)\epsilon) \rceil$, $y_1^* = (1-\psi_0)\sum_{i=1}^n \alpha_i y_i^*$ and for any x, $0 \leq \lceil x \rceil - x \leq 1$,

$$\mathbb{P}(Y_1(1) = \xi_1 | Y(0) = \xi) = \mathbb{P}(B - \mu_B \geq \xi_1 - \mu_B)$$

$$
\begin{aligned}
&\geq \mathbb{P}\left(B-\mu_B \geq \ell^*(y_1^*+b)+1-(1-\psi_0)\sum_{i=1}^n \alpha_i\xi_i\right)\\
&\geq \mathbb{P}\left(B-\mu_B \geq \ell^*\left(y_1^*-(1-\psi_0)\sum_{i=1}^n \alpha_i y_i^*\right)+\ell^* b\left(1-(1-\psi_0)\sum_{i=1}^n \alpha_i\zeta_{i-1}\right)\right.\\
&\qquad \left.+\ell^*\epsilon(1-\psi_0)\sum_{i=1}^n (i-1)\alpha_i+1\right)\\
&=\mathbb{P}\left(B-\mu_B \geq -\ell^* b(\eta-1)+\ell^*\epsilon(1-\psi_0)\sum_{i=1}^n (i-1)\alpha_i+1\right).
\end{aligned}
$$

Now, the coefficient of b in the last expression is negative due to the stability condition that $\eta > 1$ in Assumption A13. Moreover, we chose $\epsilon(\ell^*) = o(b(\ell^*))$ as $\ell^* \to \infty$, so, for sufficiently large ℓ^*,

$$
-\ell^* b(\eta-1)+\ell^*\epsilon(1-\psi_0)\sum_{i=1}^n (i-1)\alpha_i+1 \leq -\ell^*\delta := -\ell^* b(\eta-1)/2.
$$

For such ℓ^*, Hoeffding's inequality then yields

$$
\mathbb{P}(Y_1(1)=\xi_1|Y(0)=\xi) \geq 1-\exp\left(\frac{-2\ell^{*2}\delta^2}{[\ell^*(\sum_{i=1}^n \tilde{y}_i)+n]c_2^2}\right).
$$

Now, $\delta^2 = b^2(\eta-1)^2/4 = \ell^{*-2r_2}(\eta-1)^2/4$. Thus, for ℓ^* sufficiently large, we can define a constant β_0 that does not depend on ℓ^* such that

$$
\mathbb{P}(Y_1(1)=\xi_1|Y(0)=\xi) \geq 1-\exp(-\beta_0 \ell^{*\theta_3}).
$$

For the case where $z_0^* > 0$ the argument is similar, but we retain the (positive) value of z_0^* in the probability calculation, i.e., $y_1^* = (1-\psi_0)\sum_{i=1}^n \alpha_i y_i^* - z_0^*$ and $b = 0$, and then for large enough ℓ^*, ϵ is so small that $\epsilon(1-\psi_0)\sum_{i=1}^n (i-1)\alpha_i - z_0^* = -\delta < 0$, and the remainder of the argument is the same.

Using conditional independence of the various expressions above, conditional on $Y(0) = \xi$, it follows that

$$
\mathbb{P}(Y(1)=\xi|Y(0)=\xi) \geq \prod_{i=0}^{n-1}(1-e^{-\beta_i(\ell^*)^{\theta_3}}) \geq 1-\theta_1 e^{-\theta_2(\ell^*)^{\theta_3}}
$$

for appropriately defined positive constants θ_1 and θ_2 that do not depend on ℓ^*. ■

Next we show that the expected revenue received over one transition when we start in the target state is near optimal.

LEMMA EC.2. *Under A4-A14, the expected one-step revenue starting from the target state ξ is*

$$r(\xi) \geq \ell^* \rho - O(\ell^{*1-r_2}),$$

where ρ is the optimal value of the steady-state LP (4) *with $\ell^* = 1$.*

Proof: Suppose that $Y(0) = \xi$. Let $Z = Z(0)$ be the vector of sales decisions. The expected one-step revenue from selling ewes and ewe lambs is then $r(\xi) := \mathbb{E}\left[\sum_{i=0}^{n} p_i Z_i\right] = \sum_{i=0}^{n} p_i \mathbb{E} Z_i$. Now,

$$\begin{aligned}
\mathbb{E} Z_n &= (1-\psi_n)\xi_n \\
&= \ell^*(1-\psi_n) y_n^* + \ell^*(1-\psi_n)(\xi_n/\ell^* - y_n^*) \\
&= \ell^* z_n^* - O(\ell^*(b+\epsilon)).
\end{aligned}$$

Similarly, for $i = 1, 2, \ldots, n-1$, with L_i the number of lost ewes of age i at time 0,

$$\begin{aligned}
\mathbb{E} Z_i &= \mathbb{E}[\xi_i - L_i - \xi_{i+1}]_+ \\
&\geq \mathbb{E}[\xi_i - L_i - \xi_{i+1}] \\
&\geq \ell^*[(1-\psi_i)(\tilde{y}_i - (i-1)\epsilon) - (\tilde{y}_{i+1} - i\epsilon)] - 1 \\
&= \ell^*[(1-\psi_i) y_i^* - y_{i+1}^* - O(\epsilon)] - 1 \\
&= \ell^* z_i^* - O(\ell^* \epsilon).
\end{aligned}$$

Also, let B be the number of ewe lambs in year 0 that are born and not lost, so that $\mathbb{E} B = (1-\psi_0)\sum_{i=1}^{n} \alpha_i \xi_i$. Then

$$\begin{aligned}
\mathbb{E} Z_0 &= \mathbb{E}[B - \xi_1]_+ \\
&\geq \mathbb{E}[B - \xi_1] \\
&\geq \ell^* \left[(1-\psi_0) \sum_{i=1}^{n} \alpha_i (\tilde{y}_i - (i-1)\epsilon) - \tilde{y}_1 \right] - 1 \\
&= \ell^* \left[(1-\psi_0) \sum_{i=1}^{n} \alpha_i y_i^* - y_1^* - O(b+\epsilon) \right] - 1 \\
&= \ell^* z_0^* - O(\ell^*(b+\epsilon)).
\end{aligned}$$

The same argument for ewe lambs can be applied for ram lambs to establish that the expected number of non-lost ram lambs that are sold is $\mathbb{E}\mathcal{L}^R(0) \geq \ell^*(z_0^* + y_1^* - O(b+\epsilon))$. Combining the above observations with A12 yields the result. ■

To proceed further, we first need the following technical result on random walks.

LEMMA EC.3. *Suppose that* $(X_n : n \geq 1)$ *is an i.i.d. sequence with* $\mathrm{var}(X_1) < \infty$ *and* $\mathbb{E}[X_1] = 0$. *If we define the random walk* $\chi_n := X_1 + X_2 + \cdots + X_n$ *for* $n \geq 1$*, then*

$$\inf_{n\geq 1} \mathbb{P}(\chi_n \geq 0) > 0.$$

Proof of Lemma EC.3. If $\mathrm{var}(X_1) = 0$, then the result holds trivially since the infimum is 1. So assume that $\mathrm{var}(X_1) > 0$. Since $\mathbb{E}\chi_n = 0$ for all $n \geq 1$ and $\mathrm{var}(\chi_n) > 0$, it follows that $\mathbb{P}(\chi_n \geq 0) > 0$ for all $n \geq 1$. The central limit theorem proves that $\mathbb{P}(\chi_n \geq 0) \to 1/2$ as $n \to \infty$. Thus, $\mathbb{P}(\chi_n \geq 0)$ is a (strictly) positive sequence that converges to 1/2, and is therefore bounded away from 0. ■

From Lemma EC.3, it follows that there exists $q_0 > 0$ such that the (random) number of ewe-lamb births that survive their first year exceeds its mean $(1-\psi_0)\sum_{i=1}^n \alpha_i \mathbf{Y}_i(t)$ conditional on the census $\mathbf{Y}(t)$ with probability at least q_0 uniformly in $t \geq 0$, i.e., that $\mathbb{P}\left[B(t) \geq (1-\psi_0)\sum_{i=1}^n \alpha_i \mathbf{Y}_i(t) | \mathbf{Y}(t)\right] \geq q_0$ for all $t \geq 0$. (This follows by applying Lemma EC.3 to the number of births from age-i ewes for each i, and then taking the product of the lower-bounding probabilities over each age $i = 1, 2, \ldots, n$.) Moreover, a similar result holds for the number of surviving ram lambs and age-i ewes, for each $i = 1, 2, \ldots, n$. Our next result shows that, starting from a non-target state, there is a uniform (in the non-target state) upper bound on the expected time to return to the target state.

LEMMA EC.4. *Assume A4-A14. Define* $\tau := \inf\{t \geq 1 : Y(t) = \xi\}$ *as the first return time to the target state. Then, there exist nonnegative constants* κ_1, κ_2 *that do not depend on* ℓ^* *such that*

$$\sup_{y\in\mathcal{S}, y\neq\xi} \mathbb{E}[\tau | Y(0) = y] \leq \kappa_1 (\ell^*)^{\kappa_2}.$$

Proof: We use the observation that if, each year, there are sufficient births and few losses, then the population will quickly build back up to the target state ξ.

More precisely, suppose $Y(0) = y = (y_1, y_2, \ldots, y_n) \in \mathcal{S} \triangleq (y \in \mathbb{Z}_+^n : 0 \leq y \leq \xi, y \neq 0)$ with $y \neq \xi$. (Previously we used the notation y to reflect *relative* counts in the different age levels, for an overall population of size 1. For this result the interpretation differs; now the y_i variables represent unscaled counts.) Suppose first that $y_* = \min_{i=1}^n y_i = 0$. Since $y \neq 0$, there is at least one ewe in the flock. Using Lemma EC.3 we can assume a uniform lower bound, q_1 say, on the probability that the number of surviving ewes (ewes that are not "lost") exceeds its mean, for each age $1, 2, \ldots, n$.

Assume this happens for each successive transition $t = 0, 1, \ldots, n-1$. Similarly, assume that the number of non-lost lambs (lambs that are born and are not lost in their first year) exceeds its conditional mean $(1-\psi_0)\sum_{i=1}^{n} \alpha_i Y_i(t) > 0$ for $t = 0, 1, \ldots, n-1$. The conditional probabilities of these events can be bounded below by the probability $q_0 > 0$, again using Lemma EC.3. Assuming these events happen in all of the first n transitions, it follows that after n transitions, all components of the new state, $Y(n)$, are at least equal to 1. The unconditional probability of such an event is at least $(q_1^n q_0)^n$. Recall that we also sell as necessary to ensure that the ith coordinate of $Y(n) \le \xi_1$, though sales are unlikely to occur until the population recovers sufficiently.

Now we iterate. Starting from $Y(n) \ge (\zeta_0, \zeta_1, \ldots, \zeta_{n-1})$, assume the number of non-lost lambs and non-lost ewes are at least their respective conditional means in n successive steps. The first such step, with some algebra, yields a state

$$Y(n+1) \ge (\eta, \zeta_1, \zeta_2, \ldots, \zeta_{n-1}).$$

The kth such step (with $1 \le k \le n$), again with some algebra, yields a state

$$Y(n+k) \ge (\eta\zeta_0, \eta\zeta_1, \ldots, \eta\zeta_{k-1}, \zeta_k, \zeta_{k+1}, \ldots, \zeta_{n-1}).$$

Thus, after $2n$ steps, $Y(2n) \ge (\eta\zeta) \wedge \xi$, where ζ is the vector $(\zeta_0, \zeta_1, \ldots, \zeta_{n-1})$, and the componentwise minimum $\wedge$ indicates any necessary sales to avoid exceeding ξ_i ewes at any age stratum i. The probability of n such successive events is again at least $(q_1^n q_0)^n$. Now the lower bound on the population has increased by a factor $\eta > 1$, due to the stability condition, A11. With each successive sequence of n steps, the lower bound on the population increases by an additional factor of η, at least until the state hits ξ. Thus, after kn successive steps where births equal or exceed their mean, $Y(kn)$ is at least $(\eta^{k-1}\zeta) \wedge \xi$.

With each successive block of n transitions where the number of non-lost lambs and non-lost ewes are at least their respective conditional means, the lower bound on the census increases by another power of η, at least until sales are required to prevent exceeding the target state ξ. Thus, in at most $K = n\lceil(\ln \xi_1)/\ln(\eta) + 1\rceil$ successive transitions where the number of non-lost lambs equals or exceeds its conditional mean and the number of ewe losses is less than or equal to their conditional means, the state returns to the target ξ. Let $q_2 \triangleq q_1^n q_0$. Each such transition has probability at least $q_2 > 0$, so that the probability of that sequence is at least

$$q_2^K \ge q_2^{n((\ln \xi_1)/\ln(\eta)+2)} = q_2^{2n} \xi_1^{n(\ln q_2)/\ln(\eta)}.$$

Using a geometric trials argument, it follows that the expected value of τ, the hitting time to state ξ is at most

$$K/q_2^K \leq \frac{n\lceil(\ln\xi_1)/\ln(\eta)+1\rceil}{q_2^{2n}}\xi_1^{-n(\ln q_2)/\ln(\eta)} \leq \kappa_0\xi_1^{\kappa_2}$$

for appropriately defined constants $\kappa_0, \kappa_2 > 0$ that do not depend on ℓ^*. (Note that $\ln(\xi_1) \leq \xi_1$, so the impact of this term can be absorbed into κ_2, the power of ξ_1.) The observation that $\xi_1 \leq \ell^*(y_1^*+b)+1 \leq \ell^*(1+b)+1$, together with the fact that $b = b(\ell^*) < 1$ means we can take the constant $\kappa_1 = \kappa_0 2^{\kappa_2}$. ■

We now build on Lemma EC.2 on the revenue when moving from the target state to the target state to show that the long-run revenue from our policy is near optimal.

Lemma EC.5. *Under A4-A14, the mean steady-state revenue from our stated policy is at least $\ell^*\rho - O((\ell^*)^{1-r_2})$, i.e., the revenue of the optimal LP solution minus a small amount.*

Proof: The mean steady-state revenue from our policy is at least equal to the mean one-step revenue starting from the target state multiplied by the steady-state probability of being in the target state. Lemma EC.2 showed that the mean one-step revenue starting from the target state is at least $\ell^*\rho - O(\ell^{*1-r_2})$.

Now we use the fact that the steady-state probability of being in the target state is the reciprocal of $\mathbb{E}[\tau|Y(0)=\xi]$, the mean return time to the target state starting from the target state, e.g., see Resnick (2013, Proposition 2.12.3). Using first step analysis, and the bounds from Lemma EC.1 and Lemma EC.4, we have that

$$\begin{aligned}\mathbb{E}[\tau|Y(0)=\xi] &\leq 1+P(Y(1)\neq\xi|Y(0)=\xi)\sup_{y\in\mathcal{S},y\neq\xi}\mathbb{E}[\tau|Y(0)=y]\\ &\leq 1+\theta_1 e^{-\theta_2\ell^{*\theta_3}}\kappa_1\ell^{*\kappa_2}.\end{aligned}$$

It immediately follows that the mean steady-state revenue is at least

$$\frac{\ell^*\rho - O(\ell^{*1-r_2})}{1+\theta_1 e^{-\theta_2\ell^{*\theta_3}}\kappa_1\ell^{*\kappa_2}} \geq \left[\ell^*\rho - O(\ell^{*1-r_2})\right]\left[1-\theta_1 e^{-\theta_2\ell^{*\theta_3}}\kappa_1\ell^{*\kappa_2}\right] = \ell^*\rho - O(\ell^{*1-r_2}).$$

■

Lemma EC.6, below, shows that the flock size constraint is met, asymptotically.

Lemma EC.6. *Under A4-A14, the mean steady-state flock size, $\nu=\nu(\ell^*)$ say, satisfies $|\nu(\ell^*) - \ell^*| = O((\ell^*)^{1-r_2})$.*

Proof: When $Y(0)=\xi$, the flock size equals $\sum_{i=1}^{n}\xi_i$ which lies between $\ell^*(1-n(n-1)\epsilon/2)$ and $\ell^*(1+bR_n)+1$. When $Y(0)\neq\xi$, $0\leq\sum_{i=1}^{n}Y_i(0)\leq\ell^*(1+nb)+1$. If q is the steady-state probability of the target state ξ, then

$$q\ell^*(1-O(\epsilon))\leq\nu(\ell^*)\leq q\ell^*(1+bR_n)+(1-q)\ell^*(1+nb)+1,$$

and since $R_n\leq n$,

$$q\ell^*(1-O(\epsilon))\leq\nu(\ell^*)\leq\ell^*(1+bn)+1.$$

The proof of Lemma EC.5 shows that $q\geq 1-\theta_4 e^{-\theta_5\ell^{*\theta_6}}$ for some positive constants θ_4,θ_5 and $\theta_6\in(0,1)$ that do not depend on ℓ^* and $\theta_4 e^{-\theta_5\ell^{*\theta_6}}=o(\ell^{*-r_1})$. Thus,

$$-O(\ell^*\epsilon)\leq\nu(\ell^*)-\ell^*\leq O(\ell^*b).$$

■

Proof of Theorem 3: This follows immediately from Lemmas EC.5 and EC.6.