A methodology to estimate the size of an advance market commitment for vaccines to reduce enteric methane

William Arnesen*

April 11, 2025

Abstract

This paper provides a novel framework for sizing advance market commitments (AMC) when innovation requires multiple risky attempts by competing firms. The design of AMCs for early-stage technologies presents unique challenges that existing models, which assume either guaranteed success or single-firm participation, fail to address. These limitations mean that the existing literature provides little guidance for perhaps the most promising use case: incentivizing technologies whose successful invention is not imminent. This paper's model incorporates the three key features absent from previous frameworks: the probability of failure for each firm's research program, the correlation between different development pathways, and the effects of market structure. This analysis reveals that optimal AMC size depends critically on: (1) technological distance, as firms must be compensated for the time lag between research expenditures and AMC payments, (2) market concentration, as inducing additional innovation attempts becomes substantially more expensive in concentrated markets, and (3) the target probability of success, with costs rising super-linearly due to correlation between attempts. This paper then applies this framework to a case study of enteric methane vaccine development, we find that a $702.1 million commitment (present value) would be necessary to produce a two-thirds chance of successful vaccine development. The model reveals two important conclusions. First, the time lag between costs incurred and AMC funds received substantially decreases the efficiency of the AMC. Second, the importance of market concentration and correlation rises substantially as the target probability of success grows.

*email: arnesen@uchicago.edu. This research benefited from discussions with Charles Brooke (Spark Climate Solutions), Dr. Robert Banks (Global Methane Hub), and Jay Wolf (former President of the Nebraska Cattlemen).

1 Introduction

The social value of many innovations often far exceeds the private incentive to invest in their development. Market failures can severely limit innovators' ability to capture returns from their investments in three key ways. First, innovations often generate large positive externalities that benefit society broadly, but cannot be monetized by the innovator (Armitage et al., 2023). Second, many innovations face significant knowledge spillovers, where successful approaches can be copied by competitors who did not incur the research and development costs (Gillingham and Sweeney, 2010). Third, when products have a single buyer (often a government), innovators face hold-up problems where buyers can negotiate prices down to marginal cost after fixed costs are sunk, deterring initial investment (Kremer et al., 2022).

Two broad policy approaches exist to address these market failures. The first, called "push" funding, provides upfront support through mechanisms like grants, R&D tax credits, and direct funding of research institutions (Kremer and Glennerster, 2004). In contrast, "pull" funding rewards successful outcomes through mechanisms such as prizes, production tax credits, and patents. While push funding directly subsidizes innovation inputs, pull funding increases the returns to successful innovation (Kremer and Glennerster, 2004).

Advance market commitments (AMCs) represent an additional example of a pull mechanism. An AMC guarantees a subsidy per unit sold for products that meet predetermined technical specifications, effectively creating a minimum market size for innovations that do not yet exist (Kremer and Glennerster, 2004). Unlike grants, firms receive payment only upon success, allowing those with private information about their capabilities to self-select into participation. By linking payments to market adoption rather than just technical achievement, AMCs also incentivize firms to develop products that meet real-world needs (Kremer and Glennerster, 2004). Many innovations that appear promising in laboratory settings fail to achieve widespread adoption because of practical limitations unforeseen by the sponsors. For example, a new climate resilient crop variety could work perfectly in controlled conditions but prove too complex for routine use, require unrealistic changes to existing practices, or come with hidden costs that deter adoption (e.g., untrained farmers using novel NERICA-3 rice in Sierra Leone actually experienced 14% lower yields than traditional rice due to their complexity at growing) ((Glennerster et al., n.d.)). Traditional push funding or pull funding based solely on technical specifications struggle to prevent such misaligned innovations, as funders cannot anticipate all relevant practical constraints ex ante. In contrast, AMCs create a built-in 'market test' - firms only receive payment when end users actually adopt their innovation. This payment structure naturally guides firms toward designing products that balance technical performance with practical usability, as their financial returns depend not just on meeting technical specifications but on creating innovations that users willingly adopt (Kremer and Glennerster, 2004).

The most prominent real-world application of an AMC demonstrates both the potential of AMCs and the limitations of previous experience. In 2009, a coalition of donors committed $1.5 billion to purchase pneumococcal vaccines that met specified technical criteria, with manufacturers receiving subsidies proportional to their supply commitments (Kremer et al., 2020). The AMC required manufacturers to cap prices at $3.50 per dose to ensure widespread accessibility in low-income countries. This structure successfully attracted multiple manufacturers and accelerated vaccine rollout across eligible countries. However, pneumococcal AMC launched when several vaccine candidates were already in late-stage clinical trials, sidestepping many of the challenges in designing AMCs for earlier-stage innovations (Kremer et al., 2020). The primary impact was accelerating distribution rather than inducing new research and development.

More recent AMC initiatives have taken different approaches. For instance, the Frontier program for carbon dioxide removal technology negotiates bespoke contracts with individual firms, with continued funding contingent on meeting cost reduction targets (Frontier Climate, 2024). While this approach maintains a results-oriented payment structure, it sacrifices the method-agnostic nature that characterize traditional AMC design. These precedents provide limited guidance for what may be the most promising application of AMCs: incentivizing the creation of technologies whose successful invention is not imminent. Early-stage innovations face substantial technical uncertainty, require multiple firms pursuing diverse approaches, and often need significant time between initial investment and marketable products. Designing AMCs for such cases requires addressing fundamental questions about optimal subsidy size, market structure, and risk allocation that existing frameworks have not systematically explored.

To date, Kremer, Levin and Snyder (2022) provide the only theoretical framework by which to size and structure a new advance market commitment. Their model distinguishes between two cases: a technologically- close case where the AMC design focuses on capacity investment as the R&D investments are already sunk, and a technologically-distant case where both R&D and capacity investments must be incentivized.

This paper builds on their work by extending the analysis of technologically-distant innovation in two crucial ways. First, while the Kremer-Levin-Snyder model treats R&D investment as yielding deterministic outcomes (i.e. if a firm invests in R&D, then they will inevitably produce the requested innovation), this paper models innovation as inherently risky, with each research program having a significant probability of failure. Second, whereas Kremer-Levin-Snyder focuses on markets with only a single participant, this analysis incorporates market structure and competition, recognizing that multiple firms pursuing parallel research programs fundamentally changes optimal AMC design.

This framework yields three policy-relevant insights that complement and extend the Kremer-Levin-Snyder findings. First, technological distance dramat-

ically increases the required incentive size, as firms must be compensated for the gap between upfront research expenditures and delayed AMC payments at rates exceeding social discount rates. Second, market concentration substantially affects costs, as inducing additional innovation attempts becomes markedly more expensive when fewer firms can participate. Third, targeting higher probabilities of success leads to super-linear cost increases, since correlation between innovation efforts requires compensating firms for an increasingly likely scenario of split rewards.

2 Theoretical Model

2.1 Simple Model

We start with the assumption that firms will initiate a marginal research attempt if and only if the expected returns from that marginal attempt exceed the expected costs. Research is inherently risky: in many domains, such as pharmaceutical research, the overwhelming majority of research attempts ends in failure (Yamaguchi et al., 2021).

The size of the necessary pull incentive cannot be determined simply by dividing the expected costs of a research attempt by the probability of success. Under the funder's utility function, having multiple successes provides no additional value beyond a single success; what matters is the probability of achieving at least one success, not the expected number of successes.

Variable Name	**Description**
p	Probability of success
θ	Target probability of success
n	Number of innovation attempts made
X	Present value of the pull incentive
$\mathbb{E}(c)$	Present value of the expected costs of initiating an innovation attempt

The θ variable acknowledges that even with substantial investments, innovation remains inherently risky - a chance exists that all attempts may fail. The model attempts to determine the present value of a pull incentive in order to induce enough attempts such that the probability of at least one success exceeds θ, building off of Acemoglu and Linn 2004's analysis that market size has a large effect on the incentive for innovation Acemoglu and Linn (2004). This section does not address the question of the precise structure of the AMC, as different AMC structures can all have the same present value to firms, and deals only

with the question of size. The analysis begins with several assumptions, most of which will later be relaxed:

1. All firms are risk-neutral and identical ex ante.

2. Firms face one choice: whether to commit a single innovation attempt. Although firms enter sequentially, all research progresses on the same timetable.

3. All innovation attempts are identical and independent.

4. If multiple firms succeed, they split the reward proportionately. For example, if one firm succeeds, that firm receives the entire incentive X, while N successful firms would each receive $\frac{N}{2}X$.

$$\textbf{Equation 1.1} \quad \mathbb{P}(succeed|entry) = \frac{1-(1-p)^n}{n}$$

If all firms are identical, the probability that any firm receives X equals the probability of at least one success divided by the number of firms. This equation matches Assumption (4)'s splitting mechanism: if k firms succeed, a risk-neutral firm treats $\frac{1}{k}$ chance of receiving all of X as equivalent to always receiving $\frac{X}{k}$.

$$\textbf{Equation 1.2} \quad n \geq \frac{\ln(1-\theta)}{\ln(1-p)}$$

To achieve the target probability of success, one must solve for n such that the probability of at least one successful attempt exceeds θ. In other words, we must find the n such that $1-(1-p)^n \geq \theta$. Rearrangement produces Equation 1.2.

$$\textbf{Equation 1.3} \quad X \geq \frac{\ln(1-\theta)}{\theta \cdot \ln(1-p)} \cdot \mathbb{E}[c]$$

One feature of this model is that the marginal cost of inducing an additional attempt grows with the stock of existing attempts as firms now face a higher probability of splitting the prize. Figure 1 shows how, as the number of attempts rises, the marginal cost of inducing an additional attempt approaches the expected cost of innovation, and Appendix A provides the proof.

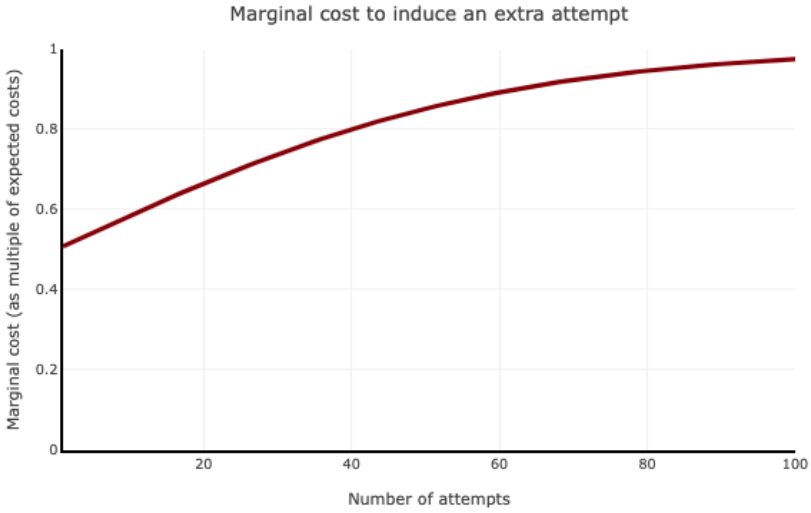

Figure 1: *The cost of inducing an additional entry approaches the cost of innovation. The higher the probability of success for each firm, the faster the convergence. In the above chart, the probability of each attempt's success is set to 5%.*

2.2 Constraints on Firm Count

2.2.1 Monopolist case

The model must now relax assumption (2), as it is unrealistic to assume firms will only make a single innovation attempt for two reasons. First, there may not be enough firms to match the number of required attempts. For example, if the probability of success per attempt is 3% and the target success probability is 75%, Equation 1.2 requires at least 46 independent attempts. While new firms could enter the market, in the short term, most markets do not have 46 firms capable of credibly achieving the target innovation. Second, firms face significant fixed costs, such as hiring specialized staff and purchasing equipment. It is unlikely that firms would incur these costs just to conduct a single trial. This analysis first considers the value equation if there was only a single firm in the market:

Equation 2.1 $$X \cdot p \cdot (1-p)^{n-1} \geq \mathbb{E}[c]$$

The monopolist requires the same number of attempts to achieve at least a θ probability of at least one successful attempt. However, the value equation from Equation 1.1 changes from the competitive market scenario. A monopolist faces not only whether to, but also how many attempts to initiate when they

do enter. The monopolist receives no additional benefit if multiple attempts succeed, and thus the funder must compensate them for the risk that the firm will have "overinvested" by having multiple successes. A marginal attempt is only worthwhile for a firm if (a) all other attempts fail, and (b) that marginal attempt succeeds.

$$\textbf{Equation 2.2} \quad X \geq \frac{1-p}{p \cdot (1-\theta)} \cdot \mathbb{E}[c]$$

We find Equation 2.2 by plugging in the value of n from Equation 1.2 into the value equation in Equation 2.1 and rearranging. The X in Equation 2.2 will always exceed the X found in Equation 1.3 for all cases where more than one attempt is required (if exactly one attempt is needed to achieve θ, then both cases are identical). In cases where a large number of attempts are required, the price premium could be substantial. For example, if the probability of success per attempt is 1% and the target probability of success is 50%, then the needed pull incentive size would be 45% larger for a monopolist than in a "one attempt per firm" approach. Section 7.2 in the Appendix shows that for all $p < \theta$, the needed pull size is higher in a monopolistic market than in a competitive market.

2.2.2 Fixed and finite firms

In practice, neither a "each firm makes one attempt" model nor a "one firm makes all attempts" model realistically maps to the real world. Instead, we introduce a variable k which represents the number of firms in the market. We keep Assumption (1) that all firms are identical, and thus if there are 40 desired attempts and four firms in the market, each will make ten attempts each.

Variable Name	Description
k	Number of firms in the market
$\mathbb{V}(X,k)$	Value of the award X given the number of firms in the market

The model follows the same logic as above, that a firm will proceed to add another attempt if and only if the marginal value they receive from that additional attempt exceeds the marginal cost of making that attempt. Due to the homogeneity assumption, each firm will initiate $\frac{n}{k}$ attempts. Future sections address the likely scenario that the desired attempts n will not be perfectly divisible by k.

$$\textbf{Equation 3.1} \quad \mathbb{V}(X,k) \cdot p \cdot (1-p)^{\frac{n}{k}-1} \geq \mathbb{E}[c]$$

Equation 3.2 $$\mathbb{V}(X,k) = \sum_{i=0}^{k-1} \frac{1}{i+1} \cdot \mathbb{P}(\text{successes} = i) \cdot X$$

Equation 3.3 $\mathbb{P}(\text{number of successful firms} = i) =$

$$\binom{k-1}{i} \cdot \left(p\right)^{\frac{n}{k}}{}^{i} \cdot \left((1-p)^{\frac{n}{k}}\right)^{k-1-i}$$

Equation 3.2 expresses that the value of having a successful attempt depends on the number of other firms that are also successful.The term $\frac{1}{i+1}$ represents the share of X received by the firm given i other successful firms. Equation 3.3 expands the binomial formula where the number of attempts per firm is $\frac{n}{k}$, equivalent to calculating the exact number of heads when one flips a weighted coin $k-1$ times where the probability of any given flip resolving as heads is $(1-p)^{\frac{n}{k}}$. Computer simulation can solve for the value of X that satisfies the above equations. In cases where the number of attempts is not divisible by the number of firms, the model applies proportionate splitting. For example, if there are 27 desired attempts and only five firms, three firms will conduct five attempts each, while two firms will conduct six.

In practice, this approach results in a higher needed pull size than the "one attempt per firm" model, but a lower pull size than the monopolist case. Intuitively, the more attempts per firm, the larger the price premium grows. For example, if the probability of success per attempt was 6% and the target probability was 50%, having only four firms results in a pull size 9.5% larger than if there were enough firms to have one attempt per firm, but 25% cheaper than if there were only a single firm. However, if the target probability was 66.67% instead, a four firm world requires a 14% premium over the one-firm-per-attempt, but a 36.3% discount from a monopolist situation.

2.3 Correlation between Attempts

The above model underestimates the necessary pull size by assuming that all attempts are independent. This assumption leads to the unrealistic result that as the number of attempts grows, the probability of at least one success approaches 1. In reality, attempts are correlated, often sharing underlying approaches, personnel, or other factors that make them dissimilar to independent coin flips.

The model introduces two sources of correlation. First, the innovation itself may be impossible, or at least infeasible within a reasonable time frame. For example, medieval scientists could have made millions of attempts to develop mRNA vaccines, yet all would fail. Second, a firm's chosen approach may be misguided. Even if a company makes 30 attempts, if they all depend on a faulty

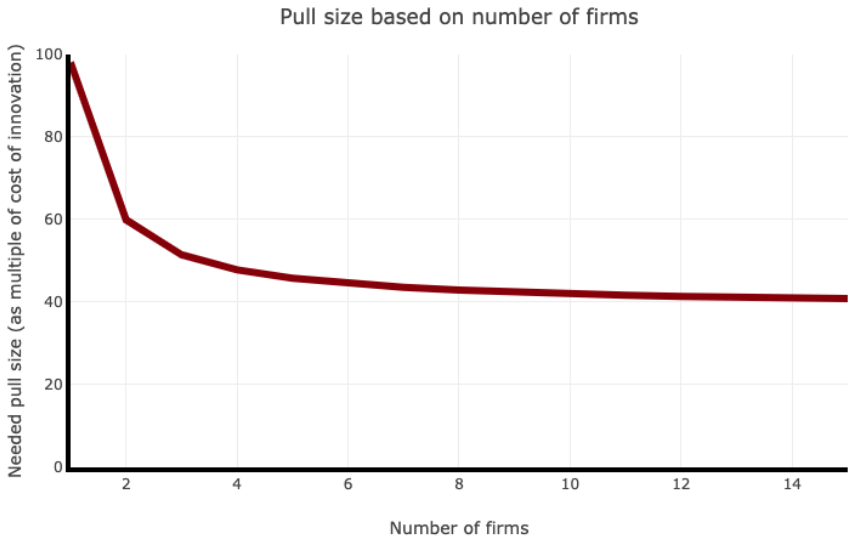

Figure 2: *The fewer the firms, the higher the needed pull size, but this value converges quickly. In the above chart, the probability of success is set to 5% and the target probability of success to 80%*

indicator, none will succeed. Thus, all attempts are correlated, and those within the same firm are even more so.

Variable Name	**Description**
η	Global possibility parameter
γ	Within-firm possibility parameter
m	Number of attempts taken by each firm

To begin, set aside the global parameter and focus only on the within-firm possibility parameter.

Equation 4.1 $\mathbb{P}_{firm}(success \geq 1) = \gamma \cdot (1 - (1 - p)^m)$

The probability that all of a firm's attempts fail is the sum of the probability that the approach is fundamentally infeasible $(1 - \gamma)$ and the probability that the approach is feasible (γ) but all of the attempts fail regardless $((1 - p)^m)$. Putting together, one gets $\mathbb{P}_{firm}(success \geq 1) = 1 - [(1 - \gamma) + \gamma \cdot (1 - p)^m]$, which simplifies to Equation 4.1.

Now to introduce the global possibility parameter η.

Equation 4.2 $\mathbb{P}_{all}(success \geq 1) = \eta \cdot (1 - [1 - \gamma \cdot (1 - (1 - p)^m)]^k) \geq \theta$

The goal is to get the probability of success to exceed the target threshold. With the new global possibility parameter, the probability that at least one firm succeeds is the probability that any given firm succeeds conditional on the innovation being possible (from Equation 4.1) multiplied by the possibility η that the innovation is at all possible.

Equation 4.3 $m \geq \dfrac{\ln(1 - \frac{1-(1-\frac{\theta}{\eta})^{\frac{1}{k}}}{\gamma})}{\ln(1-p)}$

Algebraic rearrangement of Equation 4.2 produces Equation 4.3. There is no simplified form for the necessary pull size. Instead, one can insert the value of m into Equation 3.3, replacing the $\frac{n}{k}$ term that represents attempts per firm with the value for m found in Equation 4.3.

In practice, introducing correlation substantially increases the needed pull size as well, as he number of attempts needed to reach the same probability of success increases dramatically. Subsequent sensitivity sections underscore that small changes in this value can also increase needed pull costs.

2.4 Incorporating subsidy pass-through

Our existing method estimating the appropriate size of an advance market commitment (AMC) assumes that the innovative firm receives the entirety of the subsidy pool, with no effect on the price of the novel good sold. For example, this model assumes that a three dollar per-unit subsidy would increase the profit per good sold by precisely three dollars.

This assumption is transparently erroneous. Absent binding supply constraints, increasing the profitability per good sold will induce firms to produce more. Assuming a downward-sloping demand curve, greater supply will in turn induce lower prices. By driving prices lower, the firm captures only part of the subsidy, while consumers capture the rest. A three dollar per-unit subsidy will not increase profit per good sold by three dollars, but instead some unknown value less than three dollars. We consider the transmission of the subsidy to the consumer the problem of subsidy passthrough.

In extreme cases, the passthrough could entirely eliminate the value of the subsidy. Imagine a highly stylized market for an undifferentiated good with two identical producers with a constant cost of production of two dollars per

unit. These producers are engaging in Bertrand competition, where their sole decision is to determine the price of the good and then let the market determine how much quantity is demanded at that price. These producers will each sell this good for two dollars per unit, at no profit. If one producer defected and sold at some higher price (e.g. $3 a unit), their rival can steal the entire market by charging slightly less (e.g. $2.99). As a result, the dominant strategy for both is to sell at their cost of production. If the government then introduced a subsidy per unit sold, these firms will just lower prices by the precise amount equivalent to the subsidy. For instance, suppose the government created a $0.75 subsidy. If a producer keeps their price at $2 a unit, their rival could profitably undercut them by charging a slightly lower price, since their "effective" cost of production is now only $1.25 (cost of production minus the subsidy). The dominant strategy is to just fully "pass on" the subsidy to consumers. This subsidy was not useless: the lower price benefits consumers and expands access. But the subsidy does not increase firm profitability.

In normal circumstances, subsidy passthrough is a core objective of the funder. The goal of the subsidy is to lower prices for consumers, not boost profitability for the suppliers. But AMCs are different. AMCs are well-suited for cases where expected prices are too low to justify risky and expensive research and development efforts, not just cases where prices are too high. For example, in the case of socially valuable goods like carbon capture, consumers' low willingness-to-pay means that firms have little incentive to spend money on research and development (R&D). If every increase in the subsidy simply translated one-to-one into lower prices, the subsidy will not affect the expected profitability of R&D investment.

The particular outcome one wishes to avoid is a scenario where multiple firms respond to the AMC and succeed at producing a novel product eligible for AMC funding, and competition between those two firms drives the subsidy pass-through rate to 100%. In this case, estimates produce highly unstable results. If one firm succeeds, that firm captures the entire subsidy. If two (or more) firms succeed, then the subsidy is entirely passed through to consumers. Anticipating this dynamic, firms will be highly reticent to enter, decreasing the probability that any firm succeeds. A small AMC that induces only a single firm to respond might still work as before, but increasing the AMC in order to broaden participation and induce multiple attempts at innovation might become highly cost-ineffective. For the most pressing problems, AMC designers want multiple "shots on goal" and want multiple firms pursuing disparate attempts, in hopes that at least one of them eventually succeeds. This dynamic becomes untenable with high pass-through rates.

Let us start by assuming Cournot competition, where firms decide the amount of quantity to produce and then the market sets the price. This model of competition hews more closely to reality for manufactured goods where capacity (e.g. the size of the factory) must be installed ahead of time; firms cannot

costlessly and instantaneously ramp up their supply. The precise number depends on the exact slope of the supply and demand curve, but in general, the Cournot model produces a roughly 55% reduction in profit when moving from one to two firms. More generally, each of the k-firms will receive $\frac{4}{(k+1)^2}$ of the profit of a monopolist, instead of $\frac{1}{k}$.

Installing a 10-20% premium on top of the existing model would be an oversimplification, though not a terribly inaccurate one. A small AMC where the probability of multiple successful winners is very low will need little premium. A large AMC where one targets a high probability of success (and thus a high probability of duplication) will need a larger premium. In the chart below, the y-axis represents the increase in relative payments incorporating pass-through to the previous estimates, which assumed 0% pass-through.

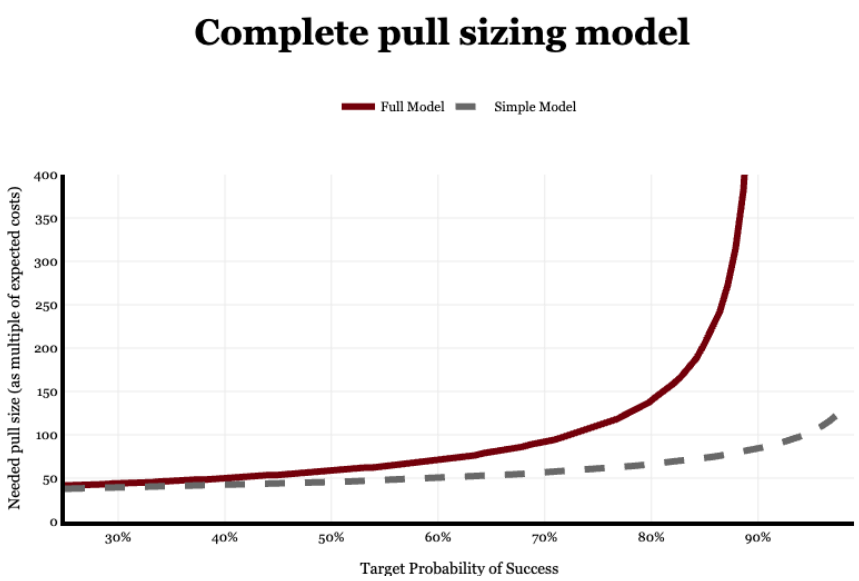

Figure 3: *The simple model (without correlation and firm constraints) diverges from the complete model as the target probability of success rises. In the above, the probability of success is set to 3%, the global possibility to 90%, the number of firms to 10, and the per-firm probability parameter to 80%.*

3 Applications for a Potential AMC to Incentivize Vaccine to Reduce Ruminant Emissions

3.1 Setting

3.1.1 Nature of the problem

Enteric methane vaccine development represents an ideal test case for optimal AMC design. The nature of the challenge combines significant social value with limited private returns, substantial technical uncertainty requiring multiple research approaches, and a concentrated market of potential innovators. These characteristics match the key features of early-stage innovation that the theoretical framework addresses.

The social value is substantial: enteric methane from ruminant animals constitutes approximately 5% of global greenhouse gas emissions when measured by warming potential (Morgavi et al., 2023). Though methane remains in the atmosphere for only about ten years, its heat-trapping potential is so potent that the EPA estimates each ton causes $1,600 in social damage – more than eight times the damage per ton of CO2 (U.S. Environmental Protection Agency, 2023). With 1.55 billion cattle worldwide (Food and Agriculture Organization of the United Nations, 2023) each generating $120-$210 in annual climate damages (U.S. Environmental Protection Agency, 2020), the aggregate social cost exceeds $200 billion annually.

Yet private incentives for innovation are effectively zero. Methane reduction provides no clear production benefits to ranchers, creating a pure externality problem. Without any private benefits, ranchers have no incentive to pay for the vaccine at any price. Even with perfect patent protection, a developer could not charge prices sufficient to recoup research and development costs because the willingness-to-pay of the end user is zero. This complete misalignment between social and private value makes the technology a representative example of when market forces alone will not induce innovation, even if the social returns would vastly exceed the costs.

The technical challenge is substantial and uncertain. A successful vaccine must trigger an immune response that reduces activity of methanogenic bacteria in the rumen without harming beneficial bacteria essential for digestion (New Zealand Agricultural Greenhouse Gas Research Centre, 2024). A diversity of potential approaches, combined with the nascent state of the science, means any single research effort faces a high probability of failure. Success likely requires multiple attempts by several firms pursuing different technical pathways.

The market structure further complicates incentive design. Veterinary vaccine development requires specialized expertise and substantial fixed costs, limiting the pool of potential innovators to a small number of large pharmaceutical firms plus several specialized biotech companies. This concentrated market

structure means each firm must make multiple attempts, increasing the importance of correlation between attempts and raising the cost of achieving a high probability of success.

These characteristics - high social value, limited private returns, technical uncertainty, and concentrated market structure - make enteric methane vaccines an ideal test case of the innovation challenges that AMCs are designed to address.

3.1.2 Existing methods to reduce ruminant emissions

Past success in reducing emissions intensity demonstrates both the potential for innovation and the limitations of relying solely on market forces. Agricultural advancements over the last century have significantly improved the emissions efficiency of cattle production, primarily by reducing the time animals spend producing methane without generating meat or milk output. Indeed, modern practices in nutrition, genetics, and herd health have enabled U.S. dairy farms to produce twice as much milk as 90 years ago with 60% fewer cows (Hristov, 2015). While today's cows produce 2.5 times more methane due to increased feed intake, emissions per kilogram of milk have dropped by 55%, from 31g in 1924 to 14g in 2014 (Hristov, 2015).

However, three factors limit the potential for further emissions reductions through existing approaches. First, increases in global demand will likely outstrip improvements in cattle productivity, resulting in net increases in emissions. The FAO projects a 12% increase in beef demand between 2024 and 2033 that will likely more than offset productivity gains (OECD/FAO, 2024). Second, agricultural management practice improvements face diminishing returns, particularly in high and middle-income countries where management practices are already optimized. Third, many strategies for improving productivity, such as increased factory farming, may raise serious environmental and ethical concerns that limit their desirability as a channel for further methane reduction (Hayek, 2022), (Verkuijl et al., 2024). While improvements in herd health and productivity can reduce emissions, significantly cutting methane levels requires investment in innovation.

Recent technical interventions further illustrate why market forces alone cannot solve this problem. Feed additives can reduce methane emissions by 4-70% (Luke and Tonsor, 2024), and in May 2024, the FDA approved the first feed ingredient clinically proven to reduce enteric methane (Elanco, 2024). However, adoption remains minimal due to misaligned incentives. These additives impose substantial recurring costs - between $70 and $105 per cow annually for Bovaer (Hanson, 2024), the only FDA-approved option - while providing no private benefits to offset these costs.

The experience with feed additives offers a crucial insight for AMC design: methane reduction typically provides no productivity benefits to offset adoption

costs. While the FAO estimates that cattle lose 12% of their energy to methane production (Food and Agriculture Organization of the United Nations, 2024), interventions like Bovaer that reduce methane do not appear to improve growth rates or feed efficiency (Hanson, 2024). This evidence suggests that even a successful innovation will face near-zero private willingness to pay, making it impossible to recoup research and development costs through market pricing alone.

3.1.3 Potential for vaccines

Preliminary evidence suggests that vaccination may offer a cheaper alternative to feed additives for reducing enteric methane. Academic studies indicate that vaccines could introduce antibodies that inhibit the production of methanogenic microbes in calves' digestive tract. Several companies, such as ArkeaBio and HelixNano, are pursuing vaccine pathways with funding from climate-oriented venture capital. However, low demand is likely to limit uptake even if a successful vaccine is developed.

In theory, vaccines are better situated than feed additives at providing a long-term solution to agricultural methane emissions. While vaccines have high upfront costs (risky R&D, capacity construction, etc.), marginal costs can be extremely low. As a result, policy mechanisms such as AMCs that increase the profitability of vaccine investment can allow firms to recoup their fixed costs and, after the AMC has expired, firms will only be left with low marginal costs of production. In contrast, feed additives involve recurring marginal costs, even as they face lower upfront costs. From a purely logistical standpoint, even if the average costs of feed additive investment was lower than that of the vaccines, inducing ranchers to continue using the feed additives requires a large, enduring subsidy program. A vaccine program would require only a notional incentive program after the end of the AMC, and needs only to cover the far lower costs of occasional injection. If the vaccine can be incorporated into other routine vaccinations (compound vaccination), then the marginal costs would approach zero.

The following sections detail the data and modeling assumptions used to estimate the size of an AMC needed to induce adequate firms to invest in R&D to construct such a vaccine.

3.2 Data & Modeling Assumptions

3.2.1 Discount Rates

The discount rate reflects how agents value money and utility today versus in the future. Different agents have varying discount factors; those with higher discount rates require greater compensation for incurring costs today on the promise of future revenues. Based on Damadoran 2024, we use the weighted average cost of capital for US pharmaceutical companies–8.05%–to represent

the discount rate for participating firms (Damadoran, 2024).

We apply a 2% discount rate for the funder and society when valuing costs and benefits. The gap between the firm and funder discount rates significantly raises expected costs: firms heavily discount future revenues, which forces the funder to raise the nominal value of the pull incentive. However, in this model, the funder does not heavily discount the expected payments they will make as much as firms. The sensitivity analysis shows how adjustments to these parameters impact the results. Generally, higher funder discount rates reduce the present value of both the pull size and expected benefits (without changing the nominal value). Higher firm discount rates, in contrast, increase the required pull size.

3.2.2 Costs Data

Parameter	**Variable**	**Value**
Firm discount rate (%)	–	8.05
Social/funder discount rate (%)	–	2
Present value of expected cost per attempt (2024 USD)	$\mathbb{E}(c)$	1,798,539
Time to develop (years)	–	8
Unconditional probability of success per attempt (%)	p	2.67
Probability the technology is possible (%)	η	75
Probability each firm's approach is feasible (%)	γ	66.67
Target probability of success (%)	θ	70

Expected cost per attempt: Expected cost per attempt ($\mathbb{E}[c]$ in the above model) is one key parameter in estimating total cost. We follow the approach and data provided by Jensen, Lund, & Fabricius 2014 to decompose innovation attempts into several phases: research & development, patent, testing, and approval (Jensen et al., 2014). We use the same cost and probability of success for each phase provided in Jensen, Lund, & Fabricius 2014 after converting into US dollars and adjusting for inflation.

Phase	**Phase Cost (2024 USD)**	**Unconditional Probability of Success (%)**	**Duration of Phase**
R&D	$2,865,030	20	3 years
Patent	$201,390	75	1 year
Testing	$305,760	20	2 years
Approval	$70,560	90	1 year

Prior to the decision to initiate an attempt, the probability that any given new attempt will succeed is thus only 2.7%. While this probability is low, if a

firm fails at any stage, they will cease incurring any further costs. We decompose the multi-year steps (R&D, testing) into annual processes:

Phase	Phase Cost (2024 USD)	Unconditional probability	Probability of reaching phase
R&D Year 1	$955,010	58%	100%
R&D Year 2	$955,010	58%	58%
R&D Year 3	$955,010	58%	33.60%
Patent	$201,390	75%	19.50%
Testing Year 1	$152,880	45%	14.60%
Testing Year 2	$152,880	45%	6.60%
Approval	$70,560	90%	3%
Final Launch	N/A	N/A	2.70%

These costs are further discounted at the 8.05% discount rate discussed above. As a result, in our model, the expected cost of initiating a new attempt is $1,798,539 with a probability of success of 2.7%.

Correlation between attempts The global possibility parameter and within-firm parameters are unobservable characteristics and vary substantially from one technology to another. As a result, they will inevitably be judgment calls, based on author research and expert interviews. The model assumes the probability that the technology is possible at 75%, but that conditional on global possibility, the probability that any given firm's approach is feasible is only 66.67%.

Conditional probability of success per attempt The model assumes that the observed success probabilities from Jensen et. al 2014 are unconditional probabilities. As a result, the probability that any given attempt will succeed conditional on a technology being feasible and a firm's approach being feasible is the quotient of the observed unconditional success probability and the product of the two possibility parameters. As a result, if a firm is on the right track (which has a probability of $75\% \cdot 66.67\% = 50\%$), that firm has a 5.3% probability of each attempt being successful. Otherwise, their probability of success per attempt is 0%.

Number of responding firms The number of possible responding firms is six, based on expert interviews and modeler judgment. As with all other parameters, the sensitivities section tests the results of the model if those assumptions change.

Target probability of success The model assumes 70%. Given that the model assumes there is only a 75% chance that the technology is possible at all, choosing an overly high probability would require inducing entry of increasingly marginal attempts, causing costs to explode upwards for increasingly slight benefits.

3.3 Modeling Assumptions

Parameter	**Value**
Cost of production ($)	2.25
Maximum share of calves vaccinated per year in the US (%)	50
AMC duration (years)	10

As the vaccine provides no private value to ranchers, the subsidy size needs to be large enough to both cover upfront and marginal costs, including cost of production, distribution and labor. The model assumes that the vaccine will be administered alongside other vaccines given to calves, so no additional distribution or labor costs are incurred.

Cost of production The cost of production was chosen at $2.25/dose, based on the cost to purchase E.Coli vaccines. E.Coli vaccines were chosen because the competitive nature of the E.Coli vaccine market suggests that sale prices for those vaccines are likely to approach their marginal cost (Lueger et al., 2012).

Uptake rate (See Figure 4) Vaccine uptake follows a logistic curve, increasing from 0% of the cattle market to 50% of new calves over 10 years. Figure 4 below shows the projected growth. Early adoption is slow, meaning that most revenues will not be realized soon after the vaccine's introduction in year 9, but closer to years 17 or 18. Since the funder applies a lower discount rate than the firm, this delay significantly increases both nominal and real costs for the funder. The eventual market size has little impact on the overall pull incentive. While a smaller market increases the per-dose subsidy, the total incentive required to justify upfront investment does not change. In fact, a smaller market may slightly reduce total costs due to lower production expenses. However, the market size has a far greater effect on the benefit side of the equation.

AMC Duration The AMC is set to last ten years after the introduction of the first vaccine. This parameter is a modeling decision to reflect both the difficulty of credibly committing funding over an indefinite period of time and the fact that the wedge between how much the funder values money after 18 years (as 70% of a dollar today) and how much a firm values the money then (as 22% of a dollar) has ballooned to such a point that it is no longer highly cost-effective to commit to funding much further.

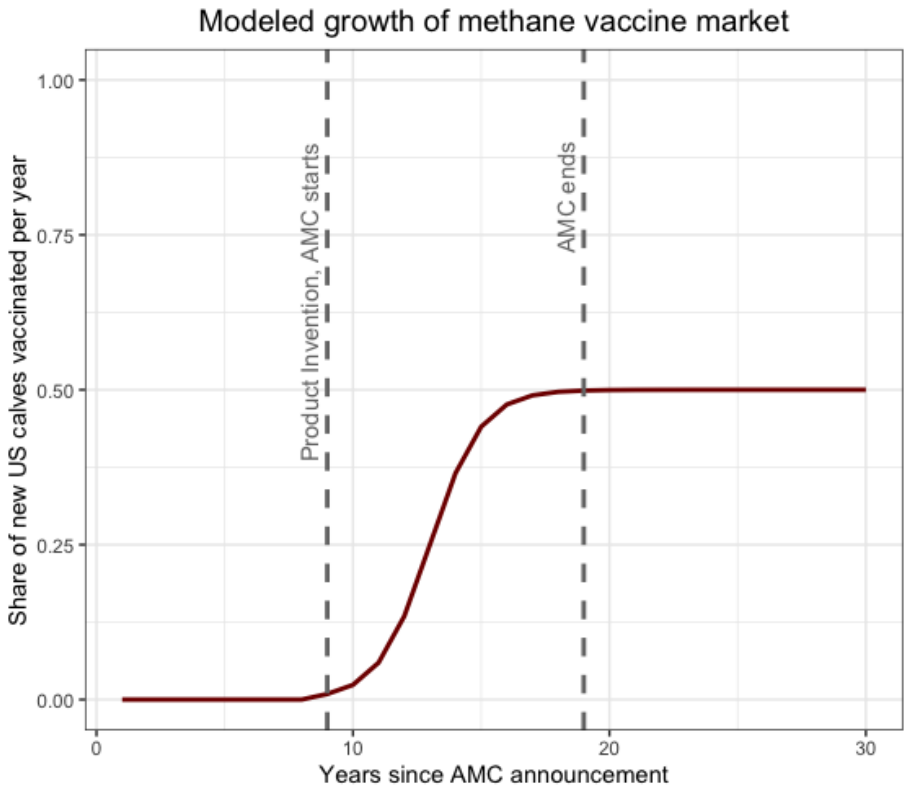

Figure 4: *We model the adoption rate as a logistic curve. Further sections test the robustness of the results to this assumption*

Subsidy per dose The cost per dose is set such that the present value of future revenues (discounted at 8.05% annually) exceeds the expected cost of development and marginal costs. A $10.33/dose subsidy, spread out over the projected 92.28 million doses during the subsidy period, yields a nominal subsidy price tag of $952.8 million, but with a present value of $702.1 million when discounted at 2% annually.

Other costs Even though the AMC ends after 10 years, any benefit-cost ratio must include social costs borne after expiry. The purchase of millions of vaccines, even if not subsidized by the funder, is still a social cost that must count in the denominator of the benefit-cost ratio. Since marginal costs pale in comparison to the upfront fixed costs and are borne distantly in the future, post-AMC non-funder social costs are only 20% of the total social costs.

Parameter	Value
Social cost of methane ($/ton)	1600
Methane per head of cattle (tons/year)	0.077
Methane reduction per dose (%)	20
Lifespan of dairy cattle (years)	5
Lifespan of beef cattle (years)	2
Maximum doses per year (million)	16.8
Share of doses to dairy cattle (%)	25
Value per dose to dairy cattle ($)	116
Value per dose to beef cattle ($)	47
Social value per vaccine ($)	64
Private value per vaccine ($)	0

3.4 Estimating Benefits

The benefits are a function of (1) the number of cattle vaccinated, (2) the kind of cattle vaccinated (beef or dairy), and (3) the timing of the vaccination.

Benefits per head of cattle: This model uses the EPA's social cost of a ton of methane at $1600/year, discounted at 2% annually. The EPA estimates that the average cow produces between 154 and 264 pounds of methane annually, which has a corresponding social cost of $123-210 annually. This model takes the most conservative assumption and assumes the minimum value ($123) per vaccinated cattle per year. The model further assumes that each vaccination reduces methane emissions by 20% and that the benefits last the lifetime of the cattle. This number is both consistent with existing conversations with leading methane vaccine companies. An AMC funder could choose a higher target threshold in exchange for a lower probability of success. The model assumes a lifespan of two years for beef cattle (75% of the herd) and five years for dairy cattle (25% of the herd), resulting in an average benefit. After discounting future benefits at a 2% rate, the present value of a vaccinated head of cattle is $116 for cattle raised for dairy, and $47 for cattle raised for beef. On expectation, each vaccination yields $64 in present value social benefit.

Period of evaluation: The benefits analysis only includes up to 30 years of benefits, of which the first eight years are consumed by vaccine development and thus have no social benefits, and the next ten contain the ramp up period in which benefits are below their peak. The analysis ends after 30 years as uncertainty grows over time and projecting the benefits too distant into the future is an exercise in folly. Vaccines may be invented in that period regardless, or alternative technologies or policies might dramatically lower the cost-effectiveness of vaccination. Indeed, this is one philosophical justification for including the 2%

annual discount rate for future benefits, even if the moral discount rate remains zero.

Other benefits: This analysis assumes there is no private benefit of methane vaccination on cattle productivity. This assumption is controversial: per the FAO, around 12% of a cow's energy is taken up with methane production so reductions in methane production could have modest productivity benefits. If so, the expected costs of an AMC would fall dramatically. If the private benefit is large enough, there may be no need for an AMC at all. Expert conversations have suggested in practice these private benefits are unlikely to materialize, and thus the model conservatively assumes them to be zero.

Total social benefits: Given the above, the model estimates total social benefits of an $654 million AMC (present value) to be $12.48 billion (present value), with a benefit cost ratio of 12.72.

3.5 Robustness Checks & Sensitivity Analysis

This section examines how variations in key model parameters and assumptions affect both the required AMC size and expected benefit-cost ratios. We begin by analyzing the fundamental model assumptions about market structure and firm behavior, then examine technical and economic parameter sensitivities, and conclude with implications for policy design.

3.5.1 No uptake after AMC expiry

One crucial assumption is that uptake will remain high even after the end of the subsidy. This assumption can only hold true if: (1) eventually a small private benefit emerges to offset the small cost of vaccination, or (2) eventually some additional, much smaller incentive scheme emerges such that it remains worthwhile for ranchers to continue vaccination even after the end of the $10.33/dose subsidy.

To test the importance of this subsidy, we can model the effects of the model if we end all vaccination after the end of the subsidy. In this case, the costs to the funders remain unchanged, the total social costs fall by 22% (since no resources need be expended during the post-subsidy period), but the total social benefits fall from $12.48 billion to $4.41 billion. The benefit cost ratio declines to 4.49.

3.5.2 Existence of a private benefit

If there is a private benefit, then the price of the AMC will fall substantially, as firms will be able to anticipate profits even outside of the AMC. If the private benefit is large enough, then no AMC may be necessary at all. If the private benefit was $2.25/dose, enough to offset the cost of production, then the AMC cost falls to $500.8 million and the benefit-cost ratio rises to 24.93. If the benefits

do allow cattle productivity to rise 2.4% (20% of the 12% energy improvements the FAO predicts) then no AMC would be necessary.

3.5.3 Parameter Changes

Parameter	Alternative Value	Current Value	New AMC Cost (PV $m)	New BCR
Baseline	-	-	702.1	12.72
Number of firms	3	6	1,281.4	8.00
Number of firms	9	6	619.8	13.88
Firm discount rate	5%	8.05%	515.6	15.7
Firm discount rate	12%	8.05%	1,073.2	9.23
Social discount rate	0.00%	2.00%	952.9	14.16
Social discount rate	4.00%	2.00%	521.1	11.47
Time to develop	4 years	8 years	601.7	16.15
Time to develop	12 years	8 years	832.8	9.14
Labor cost	$7.25/dose	0	1,195.1	5.25
Max uptake	25%	50.00%	625.6	8.16
Max uptake	100%	50.00%	855.1	17.65
Target probability	50.00%	66.67%	447.0	17.18
*Global possibility**	90.00%	75.00%	526.6	15.48
*Within-firm possibility**	50%	66.67%	908.2	10.51
*Within-firm possibility**	90%	66.67%	612.1	14.00
Unconditional probability	1.00%	2.67%	1,655.6	6.45
Unconditional probability	5.00%	2.67%	436.8	17.42
Share methane reduced	10.00%	20.00%	702.1	6.36
Share methane reduced	30.00%	20.00%	702.1	19.09
Dose wears off after 1 year	NA	NA	946.2	7.21
Years considered	20	30	702.1	7.82
Years considered	40	30	702.1	15.3
*AMC duration**	5	10	448.4	13.98
*AMC duration**	15	10	904.9	11.78

The above table reports the sensitivity to the model to changes in different parameters.

Number of firms The number of firms significantly impacts the pull incentive's effectiveness. As the number of firms decreases, each firm must make more attempts. This concentration lowers the marginal value of each additional dose because attempts by the same firm are more correlated than those by different firms, increasing the total number of attempts needed to meet the target.

Fewer firms also mean that the premium required to induce an additional attempt rises, as each additional attempt yields a smaller marginal return for the firm. Consequently, the AMC cost more than doubles, rising from $702.1 million to $1.281.4 billion, when the number of firms falls from six to three. However, the magnitude of these concentration effects depends critically on the value of the target probability of success relative to the global possibility parameter. When the target probability of success is far from the global possibility parameter, the estimate is rather insensitive to the number of firms (excluding cases of monopoly). For instance, if the target probability of success is set to 50%, then reducing the number of firms from six to three increases the needed pull size by 19% (instead of doubling). As an intuition, the number of attempts needed grows rapidly as the target probability approaches the maximum possible probability. Intuitively, designers should care most about market concentration when targeting a high probability of success.

Discount rates The higher firms' discount rates for future earnings, the more funders must commit in nominal value to compensate for the delay. Raising the annual discount rate from 8% to 12% would increase the AMC's cost from $702.1 million to over $1.073 billion. At a high discount rate of 20%, costs would exceed $2.6 billion. The social discount rate has mixed effects: a higher rate reduces the present value of future commitments (but not the nominal value) while also lowering the present value of future benefits.

Time to develop The time of development robustness check assumes that even if development time shortens from eight to four years, the nominal cost of development remains unchanged. In practice, compressing costs into earlier years increases the present value of expected development costs. However, three factors improve the benefit-cost ratio (BCR) with a shorter development time. First, earlier fund disbursements reduce the compensation needed for delay, as firms discount future revenues more steeply than funders. Second, earlier realization of benefits increases their present value. Third, shorter development time adds additional years to accrue benefits, given that the evaluation period is capped at 30 years. Overall, most of the impact of development time on the BCR comes from changes in expected benefits.

Labor costs A core assumption of the model is that labor costs are $0 because the vaccines will be administered concurrently with other vaccines. If this assumption is not true, needed costs will rise commensurate with the labor costs: every additional dollar in labor costs increases the needed subsidy per dose by $1.

Uptake Uptake has a minor effect on the size of the AMC because the primary goal is to cover the fixed costs of research and development, which do not vary with uptake. The small cost differences in the table arise from the $2.25 marginal cost incurred for each additional dose. However, benefits increase in proportion to the number of cattle vaccinated. If adoption expands globally,

uptake could significantly exceed projections.

Target probability The analysis is very sensitive to changes in the target probability. The higher the target probability, the larger the needed incentive. Since each marginal attempt is both lower value and more expensive to incentivize than the previous attempt, costs rise significantly. This effect becomes especially sharp as the target probability approaches the global possibility parameter. Of note, probability of success does not appear in the benefit cost ratio, as failure both entails zero benefits but also zero costs.

Correlation factors One reason why correlation between attempts has only a limited impact on the final benefit-cost ratio is that lowering the probability that an innovation is possible increases the conditional probability of success observed in the data (such that the unconditional probability of success remains flat at the 2.67% observed in the data). Nevertheless, global correlation has a sizable effect: the closer the global possibility parameter to the target probability, the more attempts are needed and each marginal attempt becomes exponentially more expensive.

Share methane reduced Adjusting the efficacy of the vaccine affects the benefits side of the equation but not the costs. This variable is a choice variable that depends on the technical product parameter selected. In reality, the higher the target percent reduction, the lower the probability that the innovation is possible, though this sensitivity analysis keeps the global possibility parameter constant.

A key assumption is that the vaccine remains effective throughout the cattle's lifespan. If annual re-administration is required, the per-dose subsidy decreases since more doses will be administered, but total costs increase as the mechanism must cover the production of additional doses.

Years Considered: If benefits are only considered for 20 years, total costs remain unchanged, but many long-term benefits will be excluded. This scenario is realistic if (a) uptake falls to zero after the AMC expires due to a lack of enduring incentives for ranchers to continue vaccination, or (b) a different vaccine would have been developed in the absence of the AMC. Condition (a) is especially important—without ongoing incentives, ranchers have little reason to incur the $2.25 marginal cost per dose after the AMC ends.

AMC Duration The duration of the AMC has minimal impact on costs or benefits. A shorter AMC reduces costs because firms delay revenue for a shorter period, and firms must be compensated for that delay since their discount rates exceed those of the funder.

4 Discussion

This paper advances the theory of advance market commitment design by providing the first systematic framework for sizing AMCs under conditions of risky innovation and market competition. While previous frameworks like Kremer-Levin-Snyder focus on deterministic innovation outcomes, our model incorporates three critical real-world features: probabilistic success, correlated attempts, and market structure effects. This expanded framework yields several important implications for policymakers considering AMCs as innovation incentives.

4.1 Key Policy Implications

The main implication is that the time delay between the announcement of the AMC and the expected receipt of the AMC funds has a major effect on the total real and nominal size of the AMC. If the hurdle rate exceeds the funder's social discount rate, every additional year delay increases the needed real size of the pull incentive. This relationship suggests that funders need to consider three variables:

1. **The level of technological proximity**: the closer the technology, the lower the needed cost. However, the ease of targeting more proximate technologies must be balanced against the fact that, as Kremer, Levin and Snyder 2020 note, the greater the technological proximity, the greater the private information firms have about their own capabilities. As a result, funders may need to pay higher rents in a technologically close AMC.

2. **The hurdle rate**: industries and regions with higher hurdle rates requires higher pull sizes. In particular, countries with high risk premiums see particularly high hurdle rates. Companies facing a high risk of expropriation or political instability require especially high returns in order to justify spending money at-risk (Damadoran 2024).

3. **The roll-out rate**: The faster firms can translate their innovation into actual deployment, the sooner they can get paid. The sooner firms can get paid, the smaller the incentive needs to be. While the AMC does motivate firms to work faster at distribution, some innovations face greater challenges to distribution than others. For example, innovations that require extensive marketing or convincing of consumers of their desirability, or lengthy regulatory approval, all will require a larger pull size to overcome.

4.2 Why not pay use push funding?

An astute reader may notice that the ultimate price tag far exceeds the cost of simply push funding (funding upfront) all of the necessary attempts. For example, suppose (without consideration of correlation and firm constraints),

one wished to achieve a 70% probability of a successful innovation when each attempt has a 5% probability of success. Using push funding, the cost would entail 23.5x the cost of a single innovation attempt. Even after setting aside the time delay component, a pull fund will need to be 56x the cost of a single innovation attempt. After adding in the time delay penalty, the cost of pull can rise to more than 3-4x the cost of push.

In many cases, the answer may still be that policymakers should prefer push. However, two factors are worth considering: first, this model assumes firms are homogenous, which suggests that planners are perfectly capable of identifying which firms to fund. But if firms are heterogenous, and that heterogeneity is partially invisible to the funder, then pull may become more efficient, as pull mechanisms induce firms with higher probabilities of success to self-select into participation. Whether push and pull is preferable partially depends on the extent to which funders are capable of identifying which firms are most capable firms. Second, this model assumes that innovation is binary–one either succeeds or one fails. But innovation is often a continuum–push funding, whereupon firms receive funding in advance, does not encourage firms to invest along the intensive margin in measures that increase their probability of success and in distribution, in optimizing for consumer-friendliness. In the case of a methane-reduction vaccine, push funding does not motivate firms to get shots in the legs of cows, and does not incentivize firms to invest in convincing regulators and ranchers to adopt the product. Even though the cost of a pull incentive exceeds the cost to incentivize an equivalent number of attempts using push funding, the pull funding will thus increase the quality of those attempts. In some cases, that consideration may be decisive.

4.3 Limitations of the current approach & Opportunities for future research

This paper is intended to provide a tractable, introductory means of estimating the size of an advance market commitment. However, it is not the final word – more research is necessary on several fronts in order to better estimate the needed size.

4.3.1 Firm heterogeneity

This model assumes all firms are identical. Homogeneity improves model tractability, but also obviates one of the primary motivations for pull: that firms differ in their capabilities and the funder cannot easily identify which firms are more capable than others, and thus it is more efficient to introduce a neutral mechanism that results in capable firms self-selecting into participation and incapable firms self-selecting out.

Extending this research to incorporate firm heterogeneity could have several possible effects on the needed pull size, depending on the modeler's assumptions

about the distribution of capabilities. In general, firm heterogeneity implies that achieving a lower probability of success may be cheaper than this model suggests, since a smaller pull size will only attract the most capable firms, who will have a higher probability of success than the average assumed in this model. However, the cost premium to induce increasingly marginal firms and increase the probability of success will be higher, as (a) the capability of marginal firms will be lower than the model-assumed average and (b) those marginal firms will face higher risks of splitting if they know they are facing more capable competitors. In short, the slope of the target probability-needed pull size curve gets steeper the more heterogenous the modeler assumes firms to be.

4.4 Incorporating alternative pathways

Allowing for multiple possible pathways would likely substantially reduce total costs. Targeting only a specific means of methane reduction is costly: one runs the risk that the specific means is either impossible or too expensive to implement. Moreover, each additional attempt made to achieve the same goal has diminishing returns, as these attempts are at some level correlated. Allowing for an alternative means of reducing methane would thus dramatically improve the efficiency of such an AMC. One possible pathway for reducing methane emissions from livestock includes allowing for selective breeding to spread the genes of those cattle with naturally low methane production. Achieving a high probability of success becomes more affordable with the introduction of a secondary pathway. Breeding has the potential to be substantially more cost-effective than vaccines, as the cost of establishing the research program may be substantially lower. Notably, constructing a breeding program involves little to no technological risk, and may be achievable through small-scale push funding (though pull mechanisms such as AMCs may be useful in encouraging uptake). de Haas et. al 2021 estimate that putting an economic weight on methane production during breeding could reduce methane emissions by 24% by 2050, compared to an increase by 13% in the absence of such a weight (de Haas et. al 2021). In general, the more agnostic the mechanism, the lower the probability that one is choosing an impossible target, and the more efficient an AMC becomes.

4.5 Sequential entry

This paper assumes that all firms enter near-simultaneously. Truly simultaneous entry would allow for an equilibrium where they may be over- or under-entry, where firms pursue a mixed strategy in which each firm has a certain probability of entry depending on the size of the pull incentive. The weakness of this approach is two-fold: first, in reality firms are capable of observing who else has entered. If there is "over-entry", more marginal firms can drop out. If there is "under-entry", more marginal firms can enter. Second, this approach is highly sensitive to the number of firms the modeler assumes exists in the market.

A more interesting alternative is truly sequential entry, and allowing firms

to make inferences based on the entry of other firms. The "global possibility parameter" is, in effect, a common value that other firms share but cannot directly observe. Each firm's private information is de facto a random signal about that shared value. Firm entry gives indications to potential competitors about that common value – if lots of other firms are entering, it is a likely signal that they have received a strong positive indication about the possibility. Modeling such behavior is beyond the scope of this paper and is a potential avenue for future research.

An alternative version of sequential entry would enable firms to enter if and only if other firms have failed. For example, a planner could size an AMC such that only one or two firms make attempts. If those attempts fail, then other firms enter. That adjustment would lower the needed size of the AMC, at the cost of increasing time. Incorporating this alternative is similarly beyond the scope of this paper and is a potential extension.

4.6 Risk aversion

This model further assumes all firms are risk neutral. Risk aversion will increase the needed size of the pull incentive, and will increase the relative desirability of push vis a vis pull. In a way, an increase in risk aversion would have the same effect as an increase in the firm hurdle rate: both constitute an increase in the needed compensation the funder must account for in order to motivate firms to make at-risk investments. Account for risk aversion would improve the accuracy of the model estimates.

4.7 Conclusion

This paper provides the first framework to estimate the size of an advance market commitment when innovation is both risky and costly. Previous frameworks assume that innovation is deterministic: if firms commit the needed resources, then the innovation will be created. This paper builds on that work by allowing for innovation efforts to fail. This exercise elucidates two key themes. First, the time lag between when research costs are incurred and the AMC payments are received has a major effect on the needed pull size. Second, the effects of market concentration and correlation between different attempts matter substantially more when the target probability of success grows. Further modeling to incorporate risk aversion, sequential entry, and heterogeneity among firms would likely improve the accuracy of this modeling.

References

Acemoglu, Daron and Joshua Linn, "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry," *The Quarterly Journal of Economics*, 2004, *119* (3), 1049–1090.

Armitage, Sarah C., Noël Bakhtian, and Adam B. Jaffe, "Innovation Market Failures and the Design of New Climate Policy Instruments," Working Paper 31622, National Bureau of Economic Research 2023.

Damadoran, Aswath, "Cost of Equity and Capital (US)," 2024.

Elanco, "Elanco Announces FDA Has Completed Review of Bovaer®, First-in-Class Methane-Reducing Feed Ingredient, for U.S. Dairy Industry," 2024.

Food and Agriculture Organization of the United Nations, "Number of cattle – FAO [dataset]," 2023. Retrieved November 25, 2024 from https://ourworldindata.org/grapher/cattle-livestock-count-heads.

_ , "Enteric Methane," 2024.

Frontier Climate, "Homepage," 2024. Accessed November 2024.

Gillingham, Kenneth and James Sweeney, "Market Failure and the Structure of Externalities," in "Harnessing Renewable Energy in Electric Power Systems: Theory, Practice, Policy," New York, NY: Routledge, 2010, pp. 69–91.

Glennerster, Rachel, Tavneet Suri, Jeannie Annan, Charles Dixon, and Frances Kimmins, "Shortening the Hungry Season: The Impact of Shorter Duration Rice (NERICA-3) in Sierra Leone – ATAI."

Hanson, Maureen, "What Can We Really Expect from Elanco's New Bovaer®?," *Dairy Herd Management*, 2024.

Hayek, Matthew N., "The infectious disease trap of animal agriculture," *Science Advances*, 2022, *8* (44).

Hristov, Alex, "Reducing the Emissions Intensity of Livestock Production–Case Studies of Success," *Global Research Alliance on Agricultural Greenhouse Gases*, 2015.

Jensen, J. D., M. Lund, and O. Fabricius, "Economic analysis of developing a Campylobacter vaccine to poultry: a real options approach," IFRO Report 227, Department of Food and Resource Economics, University of Copenhagen 2014.

Kremer, Michael and Rachel Glennerster, *Strong Medicine: Creating In-*

centives for Pharmaceutical Research on Neglected Diseases, Princeton University Press, 2004.

_ **, Jonathan Levin, and Christopher M. Snyder**, "Advance Market Commitments: Insights from Theory and Experience," *AEA Papers and Proceedings*, May 2020, *110*, 269–273.

_ **,** _ **,** **and** _ , "Designing Advance Market Commitments for New Vaccines," *Management Science*, 2022, *68* (7), 4786–4814.

Lueger, Aaron, Ted C. Schroeder, and David G. Renter, "Feedlot Costs of Vaccinating Cattle for *E. coli*," *AgManager.Info*, 2012.

Luke, Jaime R. and Glynn T. Tonsor, "The Enteric Methane Emission Conundrum: U.S. Beef Cattle Producer Adoption of Climate-Focused Technology," *Sustainable Production and Consumption*, October 2024, *50*, 364–375.

Morgavi, D. P., G. Cantalapiedra-Hijar, M. Eugène, C. Martin, P. Noziere, M. Popova, I. Ortigues-Marty, R. Muñoz-Tamayo, and E. M. Ungerfeld, "Review: Reducing Enteric Methane Emissions Improves Energy Metabolism in Livestock: Is the Tenet Right?," *Animal*, 2023, *17* (Suppl. 3), 100830.

New Zealand Agricultural Greenhouse Gas Research Centre, "Methane Vaccine," 2024. Accessed November 2024.

OECD/FAO, *OECD-FAO Agricultural Outlook 2024-2033*, Paris/Rome: OECD Publishing/FAO, 2024.

U.S. Environmental Protection Agency, "Agriculture and Aquaculture: Food for Thought," 2020.

_ , "Supplementary Material for the Regulatory Impact Analysis for the Final Rulemaking, "Standards of Performance for New, Reconstructed, and Modified Sources and Emissions Guidelines for Existing Sources: Oil and Natural Gas Sector Climate Review": Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances," Technical Report, National Center for Environmental Economics, Office of Policy, Climate Change Division, Office of Air and Radiation, Washington, DC 2023.

Verkuijl, Cleo, Jessie Smith, Jonathan M.H. Green, Rebecca E. Nordquist, Jeff Sebo, Matthew N. Hayek, and Maria José Hötzel, "Climate change, public health, and animal welfare: towards a One Health approach to reducing animal agriculture's climate footprint," *Frontiers in Animal Science*, 2024, *5*.

Yamaguchi, Shingo, Masayuki Kaneko, and Mamoru Narukawa, "Ap-

proval success rates of drug candidates based on target, action, modality, application, and their combinations," *Clinical and Translational Science*, 2021, *14* (3).

5 Appendix A: Proofs

5.1 Marginal cost convergence to $\mathbb{E}[c]$

First, we want to prove that the cost of inducing an additional marginal attempt approaches $\mathbb{E}[c]$ from below.

We know that the size of the incentive (X) multiplied by the probability of winning the incentive ($\frac{1-(1-p)^n}{n}$) must exceed the expected cost of an innovation attempt ($\mathbb{E}[c]$). In other words, $X \cdot \frac{1-(1-p)^n}{n} \geq \mathbb{E}[c]$. Through re-arrangement we can find that

$$X \geq \frac{n \cdot \mathbb{E}[c]}{1-(1-p)^n}$$

We need to simply take the derivative of that expression to find $\frac{\mathrm{d}X}{\mathrm{d}n}$, which we find as

$$\frac{\mathrm{d}X}{\mathrm{d}n} = \mathbb{E}[c] \cdot \frac{1-(1-p)^n + n \cdot (1-p)^n \cdot ln(1-p)}{(1-(1-p)^n)^2}$$

We know $\mathbb{E}[c] > 0$, $p > 0$, and $p \in (0,1)$. For simplicity, let $(1-p) = q$ and thus $q \in [0,1]$. Thus we get

$$\frac{\mathrm{d}X}{\mathrm{d}n} = \mathbb{E}[c] \cdot \frac{1-q^n + n \cdot q^n \cdot ln(q)}{(1-q^n)^2}$$

As n approaches infinity, q^n approaches 0. Likewise, the denominator approaches 1. The $n \cdot q^n \cdot ln(q)$ term similarly approaches 0, as the exponential q^n decays faster than the linear increase in n. Thus we get:

$$lim_{n\to\infty}(\frac{\mathrm{d}X}{\mathrm{d}n}) = \mathbb{E}[c] \cdot \frac{1-0+0}{1} = \mathbb{E}[c]$$

Next we need establish that the limit approaches $\mathbb{E}[c]$ from below. We can simply prove that the first derivative is always positive. We can easily observe that the denominator is always positive, so now we just need to prove that $1-q^n+n \cdot q^n \cdot ln(q) > 0 \ \forall q \in (0,1)$ & $n \geq 1$

We can factor out q^n and get $1-q^n \cdot (1-n \cdot \ln(q))$. Similar to above, the q^n converges to zero (polynomially) faster than the term $n \cdot \ln(q)$ becomes negative (linearly) and thus the term $q^n \cdot (1-n \cdot \ln(q)) < 1$. As a result, both the numerator and denominator are always positive on the domain $q \in (0,1)$ and $n \geq 1$. If the derivative is always positive, then the function is monotonically increasing and thus must be approaching $\mathbb{E}[c]$ from below.

5.2 Monopoly case always more expensive than competitive market

Above, we established that the needed pull size (for independent tries) for a monopolist $X \geq \frac{1-p}{p \cdot (1-\theta)}$ and the needed pull size for the competitive market is $\frac{\ln(1-\theta)}{\theta \cdot \ln(1-p)}$. We want to prove

$$\frac{1-p}{p \cdot (1-\theta)} > \frac{\ln(1-\theta)}{\theta \cdot ln(1-p)} \quad \forall p, \theta \in (0,1) \quad \& \quad p < \theta$$

We can start by re-arranging, remembering to flip the parity since $\ln(1-p)$ and $\ln(1-\theta)$ are negative. We now need to prove that

$$\frac{1-\theta}{\theta} \cdot \ln(1-\theta) > \frac{1-p}{p} \cdot ln(1-p)$$

We note that both sides of this equation take the identical form of $f(x) = \frac{1-x}{x} \cdot \ln(1-x)$. Thus we can re-write the needed equation as $f(\theta) > f(p)$. We now need only to prove that $f(x)$ is increasing over the domain, since $\theta > p$. To do so, we take the first derivative

$$f'(x) = -\frac{x + \ln(1-x)}{x^2}$$

The denominator x^2 is always positive. We observe that $x + \ln(1-x)$ is always negative: as proof, let $g(x) = x + \ln(1-x)$. $g(0) = 0$. $g'(x) = -\frac{x}{1-x} < 0$. Since $g(0) = 0$ and $g(x)$ is always decreasing, that means all values of $g(x)$ on the domain are negative and thus $x + \ln(1-x)$ is always negative. Since $f'(x) = -\frac{x+\ln(1-x)}{x^2}$ is now a negative multiplied by a negative, we can say that $f'(x)$ is always positive. If so, f is always increasing on the domain, and thus the pull size for the monopolist is always more expensive than the pull size for the competitive market.

In the case of $p \geq \theta$, the two scenarios are identical because only one attempt (and thus one firm) would ever be necessary.