

Hard-Pity Summon Calculations

Forward Calculation: “I Have Summons”

Gamba Tool

1 Purpose of the calculation

The forward calculation answers the following question:

Given a fixed number of summons and my current pity position, what is the probability of obtaining each possible number of target rewards?

The target reward may be an S-Level Hero, CH Hero, or another reward governed by a base probability and hard-pity rule. The exact banner determines the numerical values and any additional interactions, but the primary hard-pity calculation follows the same general structure.

Let the base probability of obtaining the target reward on an ordinary summon be p , and define the corresponding failure probability as:

$$q = 1 - p$$

Let H be the hard-pity length. A fresh pity cycle therefore contains at most H summons before the target reward is forced. The calculator treats the pity system as a discrete state-dependent stochastic process. It does not replace it with a binomial distribution, Gaussian approximation, or the banner’s long-run consolidated rate.

2 One hard-pity cycle

Let T denote the summon on which the next target reward is obtained, starting from fresh pity. For any summon before the hard-pity endpoint, $1 \leq t < H$, the target must fail on the previous $t - 1$ summons and then succeed on summon t :

$$P(T = t) = q^{t-1}p$$

At summon H , all probability that survives the previous $H - 1$ failures is concentrated into the forced pity result:

$$P(T = H) = q^{H-1}$$

The distribution is therefore a truncated geometric distribution with the remaining probability mass collected at the hard-pity endpoint. This hard pity creates a real discrete probability spike rather than the smooth tail one might intuitively expect (I certainly did initially) from a Gaussian distribution. We can easily check our assumptions by noting that the probabilities across the complete cycle sum to unity:

$$\sum_{t=1}^{H-1} q^{t-1}p + q^{H-1} = 1$$

3 Expected waiting time and long-run rate

One more thing before moving to how the numbers displayed on this page are computed. The expected number of summons required to obtain one target reward from fresh pity can be written using the survival probability:

$$E[T] = \sum_{t=1}^H P(T \geq t)$$

Since surviving until summon t requires $t - 1$ consecutive ordinary failures:

$$P(T \geq t) = q^{t-1}$$

Therefore we have:

$$E[T] = \sum_{j=0}^{H-1} q^j = \frac{1 - q^H}{1 - q} = \frac{1 - (1 - p)^H}{p}$$

The corresponding long-run target-reward rate implied by the explicit base-rate-plus-pity model is:

$$r_{\text{long}} = \frac{1}{E[T]} = \frac{p}{1 - (1 - p)^H}$$

This rate will be useful for intuition and for checking the model against a banner's displayed consolidated rate. However it does not generally equal the exact expectation for a finite summon session. A finite session begins at a specific pity state and often ends partway through another pity cycle, so its boundary conditions matter; starting one summon from hard pity is therefore not equivalent to starting from fresh pity with the same number of summons. The forward calculator accounts for this by retaining the current pity position as part of the probability state rather than replacing the process with a single effective success rate.

4 Exact finite-session state

Let N be the number of summons available, and let r_0 denote the initial pity position supplied by the user, and let r denote the pity position at an arbitrary later point in the calculation. The pity convention used by Gamba Tool is:

$$r = 1 \implies \text{the next summon is guaranteed}$$

Fresh pity is therefore:

$$r = H$$

Define the probability state:

$$P_n(k, r)$$

This represents the probability that after exactly n summons:

- k target rewards have been obtained
- r summons remain until the next guaranteed target reward

The initial condition is:

$$P_0(0, r_0) = 1$$

Every other state initially has probability zero.

5 State propagation

For a state with $r > 1$, the next summon is not yet forced. The state therefore branches into success and failure outcomes:

$$\begin{aligned} P_{n+1}(k+1, H) &+= p P_n(k, r) \\ P_{n+1}(k, r-1) &+= (1-p) P_n(k, r) \end{aligned}$$

A natural success increases the target count by one and resets the pity counter to H . A failure leaves the target count unchanged and decreases the remaining pity by one. When $r = 1$, the next summon is forced. There is no ordinary failure branch:

$$P_{n+1}(k+1, H) += P_n(k, 1)$$

The calculator repeats these transitions until all N summons have been propagated.

The resulting probability of obtaining exactly k target rewards is found by summing over every possible final pity position:

$$P(K_N = k) = \sum_{r=1}^H P_N(k, r)$$

This probability mass function is the underlying distribution shown by the forward-calculation histogram.

6 Guaranteed minimum

Hard pity places a strict lower bound on the number of target rewards that can be obtained. If the available summons end before the current pity guarantee is reached:

$$N < r_0$$

Then hard pity alone guarantees:

$$G(N, r_0) = 0$$

If the first guarantee is reached, one reward is guaranteed after r_0 summons and another is guaranteed after every additional block of H summons:

$$G(N, r_0) = 1 + \left\lfloor \frac{N - r_0}{H} \right\rfloor \quad N \geq r_0$$

Equivalently, the complete definition is:

$$G(N, r_0) = \begin{cases} 0 & N < r_0 \\ 1 + \left\lfloor \frac{N - r_0}{H} \right\rfloor & N \geq r_0 \end{cases}$$

This number depends only on the hard-pity structure and starting pity state. Natural early successes can improve the outcome but can never reduce it below this minimum.

7 Expected number of rewards

Once the exact probability mass function $P(K_N = k)$ is known, the expected number of target rewards is:

$$E[K_N] = \sum_k k P(K_N = k)$$

The expectation is the probability-weighted average across repeated identical summon sessions. It does not need to be an integer. An expectation of 10.3, for example, does not mean that any player can obtain 10.3 heroes. It means that the average number obtained across a sufficiently large number of identical sessions approaches 10.3.

8 Median outcome

The median provides a central integer outcome rather than an average. The Gamba Tool defines the median as the smallest integer m for which the cumulative probability reaches at least one half:

$$\text{Median}(K_N) = \min \left\{ m : \sum_{k \leq m} P(K_N = k) \geq \frac{1}{2} \right\}$$

Equivalently:

$$P(K_N \leq m) \geq \frac{1}{2}$$

The median and expectation are often close, but there is no requirement that they be equal.

9 The probability histogram

Each bar in the forward-calculation histogram represents:

$$P(K_N = k)$$

The horizontal axis is the number of target rewards obtained, while the vertical information corresponds to the exact probability of that outcome.

10 Multiple interacting pity systems

When pity systems interact, the process must retain the joint state:

$$x = (k_P, k_S, r_P, r_S), \quad \mathcal{T}_{\text{banner}}(x \rightarrow x') = P(X_{n+1} = x' \mid X_n = x)$$

Here k_P, k_S count primary and secondary rewards and r_P, r_S are their pity counters. The banner-specific one-summon transition kernel $\mathcal{T}_{\text{banner}}$ assigns each possible reward its probability and updates the reward counts, pity counters, and any additional banner state. It propagates the joint distribution as follows:

$$Q_{n+1}(x') = \sum_x Q_n(x) \mathcal{T}_{\text{banner}}(x \rightarrow x')$$

10.1 All-Hero Recruit

All-Hero uses S-Level pity 60 and A-Level pity 10, with ordinary probabilities $p_S = 0.0072$ and $p_A = 0.2068$. Natural S and A results reset their respective counters. If an S result and the A guarantee compete, the higher-rarity S result satisfies the A guarantee and resets its counter; the kernel's branches are therefore:

$$\left. \begin{array}{l} r_P = 1 \\ r_P > 1, r_S = 1 \\ r_P > 1, r_S > 1 \end{array} \right| \begin{array}{l} \text{forced S : 1} \\ \text{S : } p_S, \quad \text{forced A : } 1 - p_S \\ \text{S : } p_S, \quad \text{A : } p_A, \quad \text{neither : } 1 - p_S - p_A \end{array}$$

10.2 Rate-Up Recruit

Rate-Up uses the same coupled S/A transition structure with S pity 40, A pity 10, $p_S = 0.0097$, and $p_A = 0.03$.

10.3 Epic Recruit

Epic uses the same coupled transition structure with S pity 30, A pity 10, $p_S = 0.0343$, and $p_A = 0.1775$.

10.4 Stargazer Station

Stargaze requires one additional state component:

$$x = (k_{\text{CH}}, r_{\text{CH}}, d, \dots), \quad d = 0 \text{ when no deferred reward is waiting}$$

CH pity is 40, and every natural or pity CH resets it; summon 40 is an automatic CH. If that forced CH displaces a reward other than Omni-Acorn $\times 4$ or $\times 2$, and no reward is already pending, the displaced reward becomes d and replaces the next such Omni-Acorn result. Thus this banner's kernel propagates both CH pity and the single deferred reward.

The same banner-specific kernel drives both directions. Forward propagation is:

$$Q_{n+1}(x') = \sum_x Q_n(x) \mathcal{T}_{\text{banner}}(x \rightarrow x')$$

Backward expected remaining summons satisfy:

$$E(x) = 1 + \sum_{x'} \mathcal{T}_{\text{banner}}(x \rightarrow x') E(x'), \quad E(x) = 0 \text{ at the target state}$$

The backward calculation must retain the complete pity state because histories with equal hero counts but different pity counters have different expected remaining summons.

11 Hidden pity states

Some pity counters are visible to the player, while others are not. If a pity state is observable, the calculator can initialize it directly:

$$r_0 = r_{\text{user}}$$

If a pity state is hidden, an initial condition must instead be assumed. One possible assumption is fresh pity:

$$r_0 = H$$

Another mathematically possible approach would be to assign a probability distribution over possible hidden pity positions:

$$P_0(r)$$

The chosen treatment is not a mathematical necessity; it is a modeling assumption determined by what information the game exposes and what behavior the calculator intends to represent. The Gamba Tool therefore states banner-specific hidden-pity assumptions explicitly rather than presenting them as known player state.

12 Displayed consolidated rates

For a simple hard-pity system with ordinary probability p and pity length H , the model itself predicts the long-run consolidated rate:

$$r_{\text{model}} = \frac{p}{1 - (1 - p)^H}$$

This provides a useful consistency check against a published consolidated value. Exact agreement is not guaranteed. Differences can result from rounded published rates, undocumented mechanics, additional pity interactions, or an incorrect model assumption. Therefore, the Gamba Tool therefore distinguishes between:

Quantity	Treatment
Base rate	Used directly as a model input when the banner specifies it
Hard-pity rule	Modeled explicitly as a state transition
Current visible pity	Used directly as the initial state
Hidden pity	Initialized according to a stated banner-specific assumption
Displayed consolidated rate	Used as a consistency check rather than substituted for the explicit per-summon mechanics

The calculator does not silently alter explicit base rates simply to force the derived long-run value to match a displayed consolidated percentage.

13 Summary of the forward calculation

For a single target reward governed by base probability p , hard pity H , available summons N , and starting pity r_0 , the forward calculation begins from:

$$P_0(0, r_0) = 1$$

The state is propagated through all N summons using:

$$\begin{aligned} P_{n+1}(k+1, H) &+= pP_n(k, r) & r > 1 \\ P_{n+1}(k, r-1) &+= (1-p)P_n(k, r) & r > 1 \\ P_{n+1}(k+1, H) &+= P_n(k, 1) & r = 1 \end{aligned}$$

The final probability mass function is:

$$P(K_N = k) = \sum_{r=1}^H P_N(k, r)$$

The three primary summary statistics are:

$$G(N, r_0) = \begin{cases} 0 & N < r_0 \\ 1 + \left\lfloor \frac{N - r_0}{H} \right\rfloor & N \geq r_0 \end{cases}$$

$$E[K_N] = \sum_k kP(K_N = k)$$

$$\text{Median}(K_N) = \min \left\{ m : \sum_{k \leq m} P(K_N = k) \geq \frac{1}{2} \right\}$$

These quantities and the associated histogram are calculated from the full discrete pity-state distribution rather than from a Gaussian, binomial, or Monte Carlo approximation.