Seesaw Experimentation: A/B Tests with Spillovers *

Jin Li Ye Luo Xiaowei Zhang

November 9, 2024

Abstract

This paper examines how a firm's performance can decline despite consistently implementing successful A/B test innovations—a phenomenon we term "seesaw experimentation." While these innovations improve the measured primary dimension, they create negative externalities in unmeasured secondary dimensions that exceed the gains. Using a multivariate normal distribution model, we identify the conditions for this decline and propose positive hurdle rates as a solution. Our analysis shows how to set optimal hurdle rates to best mitigate these negative externalities and provides practical guidance for experimental design by demonstrating how these rates should vary with underlying parameters.

Keywords: A/B test, experimentation, spillovers

JEL Codes: C12, D62, L25

*Contact information: Jin Li (jli1@hku.hk), Ye Luo (kurtluo@hku.hk), Faculty of Business and Economics, The University of Hong Kong, Pokfulam Road, Hong Kong SAR. Xiaowei Zhang (xiaoweiz@ust.hk), Department of Industrial Engineering and Decision Analytics, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong SAR. We thank Shan Huang and Ruohan Zhan for helpful discussions.

1 Introduction

A/B testing has emerged as a fundamental pillar of modern business strategies, with firms leveraging these experiments to drive innovation, enhance user experiences, and boost revenue growth (Thomke (2020), Koning, Hasan and Chatterji (2022), Kohavi, Tang and Xu (2020)). Tech giants like Google and Amazon now deploy thousands of such tests annually, embedding data-driven decision-making into their innovation processes.

At the heart of A/B testing lies a crucial challenge: selecting the right metric for evaluation. This decision is inherently complex because firm performance spans multiple dimensions, and innovations often generate opposing effects across these dimensions. When a firm reduces advertisements to enhance user experience, it may sacrifice revenue. When it boosts online sales, offline purchases might decline.

These *spillovers* reflect fundamental trade-offs in business performance.[1] Managing them poses a significant challenge, particularly for large organizations where the sheer volume of tests makes interdepartmental coordination impractical. Resource constraints further complicate this challenge, as firms often cannot measure all potential effects, especially long-term impacts that resist easy quantification.

This paper examines how spillover effects shape the dynamics of experimentation. These effects arise when innovations improve performance in a primary dimension—typically aligned with current strategic priorities—while generating negative impacts in secondary dimensions. As business environments evolve, so do strategic priorities: today's focus on revenue might shift to user growth tomorrow. Consequently, even as firms accumulate successful A/B tests in their primary metrics, their overall performance may stagnate or decline.

We call this phenomenon *seesaw experimentation*—where the cumulative impact of successful A/B tests fails to generate overall business improvement. Like a seesaw, gains in one dimension are counterbalanced by losses in others, creating oscillation between different performance metrics without meaningful progress.

Using a multivariate normal distribution framework, we formally demonstrate how seesaw experimentation emerges. Our analysis reveals specific parameter conditions under which a firm can consistently succeed in individual A/B tests while experiencing long-term performance decline across all dimensions.

To address this challenge, we propose a straightforward solution: implementing a positive hurdle rate for A/B test approval. This approach ensures that only innovations exceeding a specified threshold gain implementation, effectively balancing experimentation's benefits against cumulative spillover costs. We derive the optimal hurdle rate at which the gain in the

[1] We use spillover, rather than negative spillover, for expositional ease.

primary dimension equals the negative externality. We analyze how the optimal hurdle rate varies with the underlying parameters, providing practical guidance for experimental design.

The literature addressing our specific focus is relatively small. Berman et al. (2018) examine how p-hacking in A/B testing can lead to false discoveries and impair experimental improvements. Azevedo et al. (2020) analyze how distributional characteristics, particularly fat tails, influence experimentation allocation strategies. McClellan (2022) explores how agency problems affect A/B test adoption mechanisms. While these studies provide valuable insights, they primarily focus on single-dimensional outcomes. Our work departs from this literature by explicitly modeling the interconnectedness of performance across multiple dimensions, revealing how improvements in one dimension can generate spillover effects in others.

2 Model Formulation

Consider a firm whose management is concerned with two performance dimensions, denoted as u and v. In each period $t = 1, \ldots, T$, the firm identifies its strategic priority $a_t \in \{u, v\}$ between these two dimensions. We refer to this prioritized dimension as the *primary* dimension, and the other as the *secondary* dimension.

The firm then conducts an A/B test to evaluate a potential innovation using a metric aligned with the primary dimension a_t, measuring the innovation's effect in this dimension. Let u_t and v_t represent the innovation's effects on dimensions u and v, respectively. Note that only the effect on the primary dimension is measured in the A/B test.[2]

Subsequently, the firm decides whether to adopt this innovation based on the measured effect. Let $d_t \in \{0, 1\}$ denote the adoption decision in period t, where 1 indicates adoption and 0 indicates non-adoption. The decision rule can be expressed as:

$$d_t = \mathbb{I}(a_t = u, u_t > 0) + \mathbb{I}(a_t = v, v_t > 0).$$

In other words, the firm adopts the innovation if and only if it improves the firm's performance in the primary dimension a_t for the current period t.

In this paper, we examine the firm's *cumulative performance* through a multi-period

[2]This assumption reflects practical challenges: initiators of the A/B test may not have the incentive to measure the outcomes in secondary dimensions, and, in general, outcomes are often harder to measure and may rely on private information held by different stakeholders. But even if firms can perfectly measure outcomes from the secondary dimension, their decisions depend on how to weigh the relative importance of the two. Our analysis reflects the limiting case where the importance of the primary dimension (at the moment of the decision) far outweighs the secondary dimension.

innovation process driven by A/B testing. Let

$$U_T = \sum_{t=1}^{T} d_t u_t \quad \text{and} \quad V_T = \sum_{t=1}^{T} d_t v_t.$$

Here, U_T represents the cumulative performance in dimension u of the *adopted innovations* over the time horizon T, while V_T represents the same for dimension v. We define the firm's *overall performance* at time T as the sum of these cumulative performances: $U_T + V_T$.

Suppose that the market environment changes exogenously and management accurately identifies strategic priorities to adapt over time. We model this by treating the strategic priorities a_t as independent and identically distributed (i.i.d.). Furthermore, suppose that each potential innovation's effects (u_t, v_t) on both dimensions are i.i.d. across time periods, while allowing for correlation between dimensions within each period.

Formally, we make the following assumption throughout the paper:

Assumption 1. (i) $\{a_t : t \geq 1\}$ is a sequence of i.i.d. realizations of a random variable A, where $\mathbb{P}(A = u) = p_u$, $\mathbb{P}(A = v) = p_v$, and $p_u + p_v = 1$.

(ii) $\{(u_t, v_t) : t \geq 1\}$ is a sequence of i.i.d. realizations of a random vector (U, V), following a bivariate normal distribution with mean vector m and covariance matrix Σ:

$$m = \begin{pmatrix} \mu_u \\ \mu_v \end{pmatrix} \quad \text{and} \quad \Sigma = \begin{pmatrix} \sigma_u^2 & \rho\sigma_u\sigma_v \\ \rho\sigma_u\sigma_v & \sigma_v^2 \end{pmatrix}.$$

Here, σ_u^2 and σ_v^2 are the marginal variances of U and V, respectively, and ρ is their correlation.

(iii) Both μ_u and μ_v are negative.

(iv) The sequences $\{a_t : t \geq 1\}$ and $\{(u_t, v_t) : t \geq 1\}$ are independent of each other.

Under these conditions, the firm's overall performance is a sum of i.i.d. random variables:

$$U_T + V_T = \sum_{t=1}^{T} [d_t(u_t + v_t)].$$

The strong law of large numbers then implies that

$$\mathbb{P}\left(\lim_{T \to \infty} \frac{U_T + V_T}{T} = \mathbb{E}[D(U + V)] \right) = 1.$$

Definition 1. A firm exhibits the phenomenon of *seesaw experimentation* when its long-run average performance $\mathbb{E}[D(U+V)]$ is negative despite adopting innovations only when they demonstrate positive effects in their respective primary dimensions.

Assumption 1 allow us to deliver the main message of the paper—the existence of seesaw experimentation and its mitigation via a positive hurdle rate—in the simplest and most standard framework. The conditions in this assumption can be relaxed without affecting the main points of the paper. In particular, in condition (i), a_t can be correlated over time. For condition (ii), our results also apply to broader classes of distributions, including those whose values can on discrete values. Condition (iii) is not necessary in the sense that seesaw experimentation can still occur in scenarios where either μ_u or μ_v is non-negative while the other is negative, provided that $\mu_u + \mu_v < 0$. But we assume this condition to reflect the fundamental business reality that random innovations typically do not improve secondary performance dimensions. Finally, condition (iv) is also not necessary in the sense that our main results hold even if the choice of innovation can depend on arbitrary realizations of outcomes in the past. Below, we characterize conditions under which seesaw experimentation occurs and analyze how to determine an optimal hurdle rate to address this issue.

3 Symmetric Case

We first consider the scenario where the two dimensions are symmetric. In other words, while they can be correlated, U and V have the same marginal distribution: $\mu_u = \mu_v$ and $\sigma_u = \sigma_v$. We use μ and σ to represent their common values, respectively.

Using the law of iterated expectations and the independence between A and (U, V), we can express $\mathbb{E}[D(U+V)]$ as:

$$\begin{aligned}
&\mathbb{E}[D(U+V)] \\
= &\underbrace{\mathbb{P}(A=u)}_{\text{Prob. of prioritizing dim. } u} \times \underbrace{\mathbb{P}(U>0)}_{\text{Prob. of adopting inno. in dim. } u} \times \underbrace{\mathbb{E}[U+V|U>0]}_{\text{Combined perf. if adopting inno. in dim. } u} \\
+ &\underbrace{\mathbb{P}(A=v)}_{\text{Prob. of prioritizing dim. } v} \times \underbrace{\mathbb{P}(V>0)}_{\text{Prob. of adopting inno. in dim. } v} \times \underbrace{\mathbb{E}[U+V|V>0]}_{\text{Combined perf. if adopting inno. in dim. } v} .
\end{aligned}$$

Hence, a sufficient condition for $\mathbb{E}[D(U+V)]$ to be negative is that both $\mathbb{E}[U+V|U>0]$ and $\mathbb{E}[U+V|V>0]$ are negative. These quantities can be calculated explicitly under the bivariate normal assumption.

Proposition 1. *In the symmetric case,* $\mathbb{E}[D(U+V)] < 0$ *if*

$$-1 \leq \rho < 2\left(\frac{-\mu}{\sigma}\right) M\left(\frac{-\mu}{\sigma}\right) - 1,$$

where $M(\alpha) \coloneqq (1-\Phi(\alpha))/\phi(\alpha)$ *is the* Mills ratio *of the standard normal distribution, with* ϕ *and* Φ *representing the probability density function and cumulative distribution function of the standard normal distribution, respectively.*

Proposition 1 reveals a counterintuitive finding: while a firm may seemingly march from one victory to another through successful A/B tests, its overall performance may deteriorate over time. This decline stems from significant negative spillover effects on secondary dimensions that offset the gain in the primary dimension.

Proposition 1 also provides sufficient conditions for seesaw experimentation to occur. Since $\alpha M(\alpha)$ is an increasing function of α (see Lemma 1 in the Appendix), seesaw experimentation becomes more likely as the signal-to-noise ratio $\alpha = |\mu|/\sigma$ grows, whether through an increase in $|\mu|$ or a decrease in σ. The likelihood of seesaw experimentation also increases as the correlation ρ becomes more negative. Moreover, as $(2\alpha M(\alpha) - 1)$ monotonically increases from -1 to 1 for $\alpha \in [0, \infty)$, seesaw experimentation can occur even with positive correlation.

To deal with the seesaw experimentation, firms can rely on traditional remedies such as interdepartmental coordination. But the coordination cost can be astronomical due to the sheer volume of experiments and challenges in measuring the effects on the secondary dimension. Therefore, we propose a simpler solution: the implementation of a positive hurdle rate z. This approach requires that any innovation must demonstrate a benefit exceeding z in the primary dimension before adoption, effectively raising the bar for implementation without requiring complex coordination across departments.

In this case, the firm's long-run average performance becomes:

$$\begin{aligned}\mathbb{E}[D(U+V)] &= \mathbb{P}(A=u)\mathbb{P}(U>z)\mathbb{E}[U+V|U>z] + \mathbb{P}(A=v)\mathbb{P}(V>z)\mathbb{E}[U+V|V>z] \\ &\coloneqq f(z).\end{aligned}$$

Proposition 2. *In the symmetric case, if* $\rho \in (-1, 1)$*, then* $f'(0) > 0$*,* $f(z)$ *attains its maximum at* z^* *and* $f(z^*) > 0$*, where*

$$z^* = \frac{\rho - 1}{\rho + 1}\mu.$$

Proposition 2 demonstrates that implementing a positive hurdle rate is strictly beneficial for the firm. The intuition is straightforward: When the firm sets a hurdle rate slightly above zero, it experiences two effects. First, there is a small loss in the primary dimension because some marginally beneficial innovations are not implemented. Since these are innovations that barely clear a zero hurdle rate, the foregone benefits are minimal. Second, there is a significant gain in the secondary dimension because the firm avoids implementing innovations that would have negative effects (in expectation) there. By setting a positive hurdle rate, the firm screens out potential innovations that might have small positive benefits in the primary dimension but large negative effects in the secondary dimension. This makes a positive hurdle rate optimal for the firm.

The proposition also determines the optimal hurdle rate and reveals two key relationships. First, the optimal hurdle rate is proportionally related to μ (the mean value in the secondary dimension). When μ is lower, the firm sets a higher hurdle rate. This makes intuitive sense as the firm wants to be more selective when secondary outcomes are likely to be less favorable.

Second, the proposition links the optimal hurdle rate with the correlation ρ between primary and secondary outcomes. Note that when $\rho = 0$ (no correlation), the optimal hurdle rate equals $-\mu$, which is the expected negative impact in the secondary dimension. This arises because, with no correlation, $-\mu$ is the average negative secondary effect. When ρ approaches 1 (perfect positive correlation), the optimal hurdle rate approaches zero. This is because with perfect positive correlation, any innovation that is good for the primary dimension will also be good for the secondary dimension, eliminating the need for a positive hurdle. Finally, when ρ approaches -1 (perfect negative correlation), the optimal hurdle rate approaches infinity. This extreme case occurs because perfect negative correlation means that any positive outcome in the primary dimension necessarily creates an equally large negative outcome in the secondary dimension, making the overall gain from innovation small.

The optimal hurdle rate is set by a basic economic principle: it should be where the positive gain from an innovation in the primary dimension equals the negative externality on the secondary dimension. At the optimal hurdle rate, the externality is internalized. Under the symmetric bivariate normal assumption, the negative externality is $\mathbb{E}[-V|U = u] = -\rho u - (1-\rho)\mu$. As shown in Figure 1, the gain in the primary dimension exceeds the expected negative externality only if it surpasses the optimal threshold $z^* = \frac{1-\rho}{1+\rho}\mu$.

The optimal hurdle rate in Proposition 2 maximizes the firm's long-term performance. A firm that also prioritizes short-term objectives would likely set a lower hurdle rate, accepting more innovations despite their potential negative secondary effects. This reflects the classic trade-off between short-term gains and long-term losses.

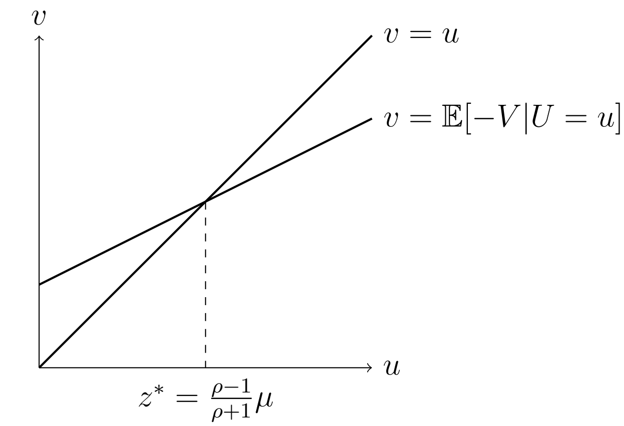

Figure 1

4 Extensions

In this section, we extend the results of Section 3 to asymmetric bivariate normal distributions and the multidimensional case.

4.1 Asymmetric Case

Consider a general bivariate normal distribution without the symmetry assumptions $\mu_u = \mu_v$ and $\sigma_u = \sigma_v$. Due to this asymmetry between dimensions, the firm sets distinct hurdle rates z_u and z_v. The firm's long-run average performance then becomes:

$$\begin{aligned}
&\mathbb{E}[D(U+V)] \\
&= \mathbb{P}(A=u)\mathbb{P}(U>z_u)\mathbb{E}[U+V|U>z_u] + \mathbb{P}(A=v)\mathbb{P}(V>z_v)\mathbb{E}[U+V|V>z_v] \\
&:= g(z_u, z_v). \qquad (1)
\end{aligned}$$

Proposition 3. *In the asymmetric case, $g(0,0) < 0$ if $-1 \leq \rho < \min(\rho_1, \rho_2)$, where*

$$\rho_1 = \left[\left(\frac{-\mu_u}{\sigma_u}\right) M\left(\frac{-\mu_u}{\sigma_u}\right)\left(1+\frac{\mu_v}{\mu_u}\right) - 1\right]\frac{\sigma_u}{\sigma_v},$$
$$\rho_2 = \left[\left(\frac{-\mu_v}{\sigma_v}\right) M\left(\frac{-\mu_v}{\sigma_v}\right)\left(1+\frac{\mu_u}{\mu_v}\right) - 1\right]\frac{\sigma_v}{\sigma_u}.$$

Proposition 4. *In the asymmetric case, if*

$$-\min\left(\frac{\sigma_v}{\sigma_u}, \frac{\sigma_u}{\sigma_v}\right) < \rho < \min\left(\frac{\mu_v/\sigma_v}{\mu_u/\sigma_u}, \frac{\mu_u/\sigma_u}{\mu_v/\sigma_v}\right),$$

then $\nabla g(0,0) > 0$, $g(z_u, z_v)$ *attains its maximum at* (z_u^*, z_v^*), *and* $g(z_u^*, z_v^*) > 0$ *where*

$$z_u^* = \frac{\rho\mu_u/\sigma_u - \mu_v/\sigma_v}{\rho/\sigma_u + 1/\sigma_v} \quad \textit{and} \quad z_v^* = \frac{\rho\mu_v/\sigma_v - \mu_u/\sigma_u}{\rho/\sigma_v + 1/\sigma_u}.$$

The conditions in Proposition 3 reduce to that in Proposition 1 when the means and variances of the two dimensions are equal. Proposition 4 derives the optimal hurdle rates for both dimensions. The optimal hurdle rate for each dimension follows the same principle as in the symmetric case: it is set to the value at which the positive gain from adopting an innovation exactly offsets its negative externality on the other dimension.

The optimal hurdle rate increases when either the mean of the secondary dimension is lower or when the correlation coefficient is lower, following the same logic as in the symmetric case since both conditions lead to greater negative externalities.

In addition, the optimal hurdle rate also depends on the mean of the primary dimension, except when the correlation coefficient ρ is zero. In the special case where $\rho = 0$, the hurdle rate equals $-\mu$ (the negative externality imposed on the other dimension). When ρ is negative, the optimal hurdle rate z_u^* decreases as μ_u increases, because a higher mean in the primary dimension reduces the negative externality, leading to a higher optimal hurdle rate.[3] Similarly, when ρ is positive, z_u^* increases in μ_u.

4.2 Multidimensional Case

Now suppose the firm measures performance across n dimensions. Each period, it identifies a dimension A as its strategic priority, evaluates a potential innovation designed for dimension A through an A/B test, and adopts the innovation if its effect on dimension A exceeds a threshold z. Let X_i denote the innovation's effect on dimension i ($i = 1, \ldots, n$), although only the effect on the primary dimension A is measured. Define the adoption decision as $D = \mathbb{I}(X_A > z)$. The long-run average of the firm's overall performance is then:

$$\mathbb{E}\left[D\sum_{i=1}^{n} X_i\right] = \sum_{i=1}^{n} \mathbb{P}(A=i)\mathbb{P}(X_j > z)\mathbb{E}\left[D\sum_{i=1}^{n} X_i \middle| X_j > z\right] := h(z). \tag{2}$$

Proposition 5. *Assume* $(X_1, \ldots, X_n)$ *follows a multivariate normal distribution with* $\mathbb{E}[X_i] = \mu < 0$ *and* $\mathrm{Var}[X_i] = \sigma^2$ *for all* i, *and* $\mathrm{Cov}[X_i, X_j] = \rho\sigma^2$ *for all* $i \neq j$. *Then,* $h(z) < 0$ *if*

$$-\frac{1}{n-1} \leq \rho < \left(\frac{n}{n-1}\right)\left(\frac{-\mu}{\sigma}\right) M\left(\frac{-\mu}{\sigma}\right) - \frac{1}{n-1}.$$

[3] Under the bivariate normal assumption, the negative externality is $\mathbb{E}[-V|U] = \rho(\sigma_v/\sigma_u)(\mu_u - U) - \mu_v$. As a result, when $\rho < 0$, an increase in μ_u reduces the negative externality.

Proposition 6. *Under the multivariate normal assumption of Proposition 5, if* $-1/(n-1) < \rho < 1$*, then* $h'(0) > 0$*,* $h(z)$ *attains its maximum at* z^**, and* $h(z^*) > 0$*, where*

$$z^* = \frac{(n-1)(\rho-1)}{1+(n-1)\rho}\mu.$$

When $n = 2$, the expressions in Propositions 5 and 6 reduce to those in our bivariate-normal analysis. In Proposition 5, the lower bound $-1/(n-1)$ is the lowest feasible correlation coefficient in a multivariate normal distribution. This indicates that seesaw experimentation emerges when A/B test outcomes across different dimensions are sufficiently negatively correlated, consistent with our earlier findings. Proposition 6 shows that the optimal hurdle rate increases in n. This reflects an intuitive relationship: more dimensions create more negative externalities, leading to a higher optimal threshold for innovation adoption.

5 Conclusion

Our paper examines A/B testing in the presence of spillover effects. Using a multivariate normal framework, we identify conditions that give rise to seesaw experimentation—where continuous adoption of successful innovations leads to performance decline. We show that this phenomenon is more likely to occur when the signal-to-noise ratio is large and correlations across dimensions are more negative. We also show how implementing appropriate hurdle rates can enhance firm performance, and derive optimal thresholds within our analytical framework. At the optimal threshold, the positive gain is equal to the negative externality it imposes.

While our model employs a multivariate normal distribution, the core mechanism we identify extends beyond this specific framework. The seesaw phenomenon and the benefits of hurdle rates persist across various outcome distributions, making our insights broadly applicable. Setting optimal hurdle rates in practice requires careful measurement of innovation outcomes across dimensions, including their correlations. For this purpose, firms can leverage their historical A/B test data, which provides a rich source of information about both direct effects and spillovers.

Our analysis demonstrates that positive hurdle rates emerge as a natural response to negative spillovers across dimensions. This principle extends symmetrically: in environments with positive spillovers, negative hurdle rates would optimally encourage innovation adoption. While the precise value of optimal hurdle rates may vary by context, our findings highlight a fundamental principle: when innovations have multi-dimensional impacts, firms should move beyond zero-threshold A/B testing. The interconnected nature of modern organizations demands a more systematic approach to experimentation—one that explicitly accounts for these cross-dimensional externalities. This insight opens new pathways for firms to design more effective experiments.

References

Azevedo, Eduardo M., Alex Deng, José Luis Montiel Olea, Justin Rao, and E. Glen Weyl. 2020. "A/B Testing with Fat Tails." *Journal of Political Economy*, 128(12): 4614–000. 2

Berman, Ron, Leonid Pekelis, Aisling Scott, and Christophe Van den Bulte. 2018. "p-Hacking in A/B Testing." *working paper.* 2

Johnson, Norman L., and Samuel Kotz. 1972. *Distribution in Statistics: Continuous Multivariate Distributions.* John Wiley & Sons. 11

Kohavi, Ron, Diane Tang, and Ya Xu. 2020. *Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing.* Cambridge University Press. 1

Koning, Rembrand, Sharique Hasan, and Aaron Chatterji. 2022. "Experimentation and Start-up Performance: Evidence from A/B Testing." *Management Science*, 68(9): 6434–6453. 1

McClellan, Andrew. 2022. "Experimentation and Approval Mechanisms." *Econometrica*, 90(5): 2215–2247. 2

Thomke, Stefan. 2020. *Experimentation Works: The Surprising Power of Business Experiments.* Harvard Business Review Press. 1

A Proofs for Bivariate Normals

Lemma 1. *$g(\alpha) \coloneqq \alpha M(\alpha)$ is strictly increasing in $\alpha \in \mathbb{R}$, $g(0) = 0$, and $\lim_{\alpha\to\infty} g(\alpha) = 1$.*

Proof of Proposition 1. Note: $g(\alpha) = \alpha(1-\Phi(\alpha))/\phi(\alpha)$ and $\phi'(\alpha) = (-\alpha)\phi(\alpha)$. Hence,

$$g'(\alpha) = \frac{1-\Phi(\alpha)+\alpha(-\phi(\alpha)) - (-\alpha\phi(\alpha))}{\phi^2(\alpha)} = \frac{1-\Phi(\alpha)}{\phi^2(\alpha)} > 0.$$

Moreover, the limit of $g(\alpha)$ can be calculated via L'Hôpital's rule. □

Lemma 2. *Let $g(z_u, z_v)$ be as defined in (1). Then,*

$$\begin{aligned} g(z_u, z_v) = p_u &\left[(\mu_u+\mu_v)\left(1-\Phi\left(\frac{z_u-\mu_u}{\sigma_u}\right)\right) + (\sigma_u+\rho\sigma_v)\,\phi\left(\frac{z_u-\mu_u}{\sigma_u}\right)\right] \\ &+ p_v\left[(\mu_u+\mu_v)\left(1-\Phi\left(\frac{z_v-\mu_v}{\sigma_v}\right)\right) + (\rho\sigma_u+\sigma_v)\,\phi\left(\frac{z_v-\mu_v}{\sigma_v}\right)\right]. \end{aligned} \tag{3}$$

Proof of Lemma 2. We first note that

$$\mathbb{E}[D(U+V)] = p_u\mathbb{E}[U+V|U>z_u]\mathbb{P}(U>z_u) + p_v\mathbb{E}[U+V|V>z_v]\mathbb{P}(V>z_v).$$

The expression (3) can be derived by using the following properties of the truncated normal distribution (Johnson and Kotz 1972, p. 112–113): for any z,

$$\mathbb{E}[U|U>z] = \mu_u + \sigma_u\lambda\left(\frac{z-\mu_u}{\sigma_u}\right) \quad \text{and} \quad \mathbb{E}[V|U>z] = \mu_v + \rho\sigma_v\lambda\left(\frac{z-\mu_u}{\sigma_u}\right),$$

where $\lambda(\alpha) \coloneqq 1/M(\alpha)$ is the inverse Mills ratio of the standard normal distribution. □

Proof of Proposition 1. The proof follows from $f(0) = g(0,0)$ and Lemma 2 through direct calculations. □

Proof of Proposition 2. Note that

$$\begin{aligned} f(z) = g(z,z) &= 2\mu\left(1-\Phi\left(\frac{z-\mu}{\sigma}\right)\right) + \sigma(1+\rho)\phi\left(\frac{z-\mu}{\sigma}\right), \\ f'(z) &= -\frac{1}{\sigma}\phi\left(\frac{z-\mu}{\sigma}\right)\left[2\mu + (1+\rho)(z-\mu)\right], \\ f''(z) &= -\frac{1}{\sigma}\phi\left(\frac{z-\mu}{\sigma}\right)\left[\left(-\frac{1}{\sigma}\right)\left(\frac{z-\mu}{\sigma}\right)(2\mu+(1+\rho)(z-\mu)) + 1 + \rho\right]. \end{aligned}$$

Hence, if $-1 < \rho < 1$, then $f'(0) > 0$ and $f'(z) = 0$ has a unique solution

$$z^* = \frac{\rho - 1}{\rho + 1}\mu;$$

moreover,

$$f''(z^*) = -\frac{1}{\sigma}\phi\left(\frac{-2\mu}{\sigma(1+\rho)}\right)(1+\rho) < 0,$$

which implies z^* the a maximizer of $f(z)$. Moreover,

$$\begin{aligned} f(z^*) &= \sigma(1+\rho)\phi\left(\frac{z^*-\mu}{\sigma}\right)\left[\frac{2\mu}{\sigma(1+\rho)}M\left(\frac{z^*-\mu}{\sigma}\right)+1\right] \\ &= \sigma(1+\rho)\phi\left(\frac{z^*-\mu}{\sigma}\right)\left[-\left(\frac{z^*-\mu}{\sigma}\right)M\left(\frac{z^*-\mu}{\sigma}\right)+1\right] > 0, \end{aligned}$$

where the positivity follows from $\alpha M(\alpha) < 1$ by Lemma 1. □

Proof of Proposition 3. By equation (3), a sufficient condition for $g(0,0) < 0$ is

$$\begin{cases} (\mu_u+\mu_v)\left(1-\Phi\left(\frac{-\mu_u}{\sigma_u}\right)\right) + (\sigma_u+\rho\sigma_v)\,\phi\left(\frac{-\mu_u}{\sigma_u}\right) < 0, \\ (\mu_u+\mu_v)\left(1-\Phi\left(\frac{-\mu_v}{\sigma_v}\right)\right) + (\rho\sigma_u+\sigma_v)\,\phi\left(\frac{-\mu_v}{\sigma_v}\right) < 0. \end{cases} \tag{4}$$

By straightforward calculations, this condition is the same as $\rho < \min(\rho_1, \rho_2)$, where

$$\begin{aligned} \rho_1 &= \left[\left(\frac{-\mu_u}{\sigma_u}\right)M\left(\frac{-\mu_u}{\sigma_u}\right)\left(1+\frac{\mu_v}{\mu_u}\right)-1\right]\frac{\sigma_u}{\sigma_v}, \\ \rho_2 &= \left[\left(\frac{-\mu_v}{\sigma_v}\right)M\left(\frac{-\mu_v}{\sigma_v}\right)\left(1+\frac{\mu_u}{\mu_v}\right)-1\right]\frac{\sigma_v}{\sigma_u}. \end{aligned}$$

□

Proof of Proposition 4. Note that

$$\frac{\partial g(z_u,z_v)}{\partial z_u} = p_u\left(-\frac{1}{\sigma_u}\right)\phi\left(\frac{z-\mu_u}{\sigma_u}\right)\left[(\mu_u+\mu_v)+(\sigma_u+\rho\sigma_v)\left(\frac{z-\mu_u}{\sigma_u}\right)\right],$$

$$\frac{\partial g(z_u,z_v)}{\partial z_v} = p_v\left(-\frac{1}{\sigma_v}\right)\phi\left(\frac{z-\mu_v}{\sigma_v}\right)\left[(\mu_u+\mu_v)+(\rho\sigma_u+\sigma_v)\left(\frac{z-\mu_v}{\sigma_v}\right)\right].$$

Hence, $\nabla g(0,0) > 0$ if and only if

$$(\mu_u + \mu_v) + (\sigma_u + \rho\sigma_v)\left(\frac{-\mu_u}{\sigma_u}\right) < 0 \quad \text{and} \quad (\mu_u + \mu_v) + (\rho\sigma_u + \sigma_v)\left(\frac{-\mu_v}{\sigma_v}\right) < 0,$$

which is equivalent to

$$\rho < \frac{\mu_v/\sigma_v}{\mu_u/\sigma_u} \quad \text{and} \quad \rho < \frac{\mu_u/\sigma_u}{\mu_v/\sigma_v}.$$

If $\rho > -\min(\sigma_v/\sigma_u, \sigma_u/\sigma_v)$, then $\nabla g(z_u, z_v) = 0$ has a unique solution

$$z_u^* = \frac{\rho\mu_u/\sigma_u - \mu_v/\sigma_v}{\rho/\sigma_u + 1/\sigma_v} \quad \text{and} \quad z_v^* = \frac{\rho\mu_v/\sigma_v - \mu_u/\sigma_u}{\rho/\sigma_v + 1/\sigma_u};$$

moreover,

$$\frac{\partial^2 g(z_u^*, z_v^*)}{\partial z_u^2} = p_u\left(-\frac{1}{\sigma_u}\right)\phi\left(\frac{-(\mu_u + \mu_v)}{\sigma_u + \rho\sigma_v}\right)\left(1 + \frac{\sigma_v}{\sigma_u}\rho\right) < 0,$$

$$\frac{\partial^2 g(z_u^*, z_v^*)}{\partial z_v^2} = p_v\left(-\frac{1}{\sigma_v}\right)\phi\left(\frac{-(\mu_u + \mu_v)}{\rho\sigma_u + \sigma_v}\right)\left(1 + \frac{\sigma_u}{\sigma_v}\rho\right) < 0,$$

$$\frac{\partial^2 g(z_u^*, z_v^*)}{\partial z_u \partial z_v} = \frac{\partial^2 g(z_u^*, z_v^*)}{\partial z_v \partial z_u} = 0.$$

Thus, $\nabla^2 g(z_u^*, z_v^*)$ is negative definite, implying (z_u^*, z_v^*) is the maximizer of $g(z_u, z_v)$. Moreover,

$$\begin{aligned}
g(z_u^*, z_v^*) &= p_u(\sigma_u + \rho\sigma_v)\phi\left(\frac{z_u^* - \mu_u}{\sigma_u}\right)\left[\left(\frac{\mu_u + \mu_v}{\sigma_u + \rho\sigma_v}\right)M\left(\frac{z_u^* - \mu_u}{\sigma_u}\right) + 1\right] \\
&\quad + p_v(\sigma_u + \rho\sigma_v)\phi\left(\frac{z_v^* - \mu_v}{\sigma_v}\right)\left[\left(\frac{\mu_u + \mu_v}{\rho\sigma_u + \sigma_v}\right)M\left(\frac{z_v^* - \mu_v}{\sigma_v}\right) + 1\right] \\
&= p_u(\sigma_u + \rho\sigma_v)\phi\left(\frac{z_u^* - \mu_u}{\sigma_u}\right)\left[-\left(\frac{z_u^* - \mu_u}{\sigma_u}\right)M\left(\frac{z_u^* - \mu_u}{\sigma_u}\right) + 1\right] \\
&\quad + p_v(\sigma_u + \rho\sigma_v)\phi\left(\frac{z_v^* - \mu_v}{\sigma_v}\right)\left[-\left(\frac{z_v^* - \mu_v}{\sigma_v}\right)M\left(\frac{z_v^* - \mu_v}{\sigma_v}\right) + 1\right] > 0,
\end{aligned}$$

where the positivity follows from Lemma 1, which implies $\alpha M(\alpha) < 1$, and the assumption that $\rho > -\min(\sigma_v/\sigma_u, \sigma_u/\sigma_v)$, which implies $\sigma_u + \rho\sigma_v > 0$ and $\rho\sigma_m u + \sigma_v > 0$. □

B Proofs for Multivariate Normals

Lemma 3. *Let $h(z)$ be as defined in (2). Then,*

$$h(z) = \mathbb{E}\left[D\sum_{i=1}^{n} X_i\right] = n\mu\left(1 - \Phi\left(\frac{z-\mu}{\sigma}\right)\right) + \sigma[1+(n-1)\rho]\phi\left(\frac{z-\mu}{\sigma}\right).$$

Proof of Lemma 3. The proof follows similar calculations to those in Lemma 2. □

Proof of Proposition 5. The proof follows from Lemma 3 through direct calculations. □

Proof of Proposition 6. Note that

$$\begin{aligned} h'(z) &= -\frac{1}{\sigma}\phi\left(\frac{z-\mu}{\sigma}\right)[n\mu + [1+(n-1)\rho](z-\mu)] \\ h''(z) &= -\frac{1}{\sigma}\phi\left(\frac{z-\mu}{\sigma}\right)\left[\left(-\frac{1}{\sigma}\right)\left(\frac{z-\mu}{\sigma}\right)(n\mu + [1+(n-1)\rho](z-\mu)) + 1 + (n-1)\rho\right]. \end{aligned}$$

Hence, if $-1/(n-1) < \rho < 1$, then $h'(z) = 0$ has a unique solution

$$z^* = \frac{(n-1)(\rho-1)}{1+(n-1)\rho}\mu;$$

moreover,

$$h''(z^*) = -\frac{1}{\sigma}\phi\left(\frac{-n\mu}{\sigma[1+(n-1)\rho]}\right)[1+(n-1)\rho] < 0,$$

which implies z^* is the maximizer of $h(z)$. Moreover,

$$\begin{aligned} h(z^*) &= \sigma[1+(n-1)\rho)]\phi\left(\frac{z^*-\mu}{\sigma}\right)\left[\frac{n\mu}{\sigma[1+(n-1)\rho]}M\left(\frac{z^*-\mu}{\sigma}\right) + 1\right] \\ &= \sigma[1+(n-1)\rho)]\phi\left(\frac{z^*-\mu}{\sigma}\right)\left[-\left(\frac{z^*-\mu}{\sigma}\right)M\left(\frac{z^*-\mu}{\sigma}\right) + 1\right] > 0, \end{aligned}$$

where the positivity follows from Lemma 1 and the assumption that $\rho > -1/(n-1)$. □