Internet Appendix

Dynamic Resource Allocation with Hidden Volatility

by Felix Zhiyu Feng and Mark M. Westerfield

Journal of Financial Economics, Forthcoming

February, 2020

B Model Robustness

B.1 Private Benefits with Volatility

We can extend the model to consider situations in which private benefits generate cash flow volatility. In the model introduced in Section 2, resources can be used to generate a flow of private benefits to the agent equal to $\lambda(K_t - \hat{K}_t)dt$. We now additionally assume that these private benefits create cash flow volatility, so that

$$\Sigma_t = f(\hat{K}_t)\hat{\sigma}_t + \lambda(K_t - \hat{K}_t)\epsilon, \tag{B.1}$$

where ϵ is the cash flow volatility per unit of private benefits. The arguments of Section 3 follow, with

$$\beta(\sigma, K) = \frac{\lambda}{f'(K)\left(\mu(\sigma) - \sigma\mu'(\sigma)\right) + \lambda\epsilon\mu'(\sigma)}. \tag{B.2}$$

If $\lambda\epsilon - \sigma f'(K) < 0$, then $\frac{\partial}{\partial\sigma}\beta(\sigma, K) < 0$, as in the baseline model.

In contrast, if $\lambda\epsilon - \sigma f'(K) > 0$, then $\frac{\partial}{\partial\sigma}\beta(\sigma, K) > 0$. This is interpretable: if re-allocating resources from productive to unproductive use greatly increases volatility, then a matching reduction must be made in productive cash flow volatility. Incentives are used to push the other way – to encourage more volatility in the productive assets. However, it is also the case that stronger incentives are used to produce less *efficient* risk taking. Recall from Section 3 that $\mu(\sigma)$ is hump shaped, and $\frac{\mu(\sigma)}{\sigma}$ is decreasing in the relevant region. If $\beta_\sigma > 0$, then stronger incentives and high volatility are associated with a lower return-to-risk ratio.

Similarly, if $\beta_\sigma > 0$ everywhere, then the arguments in Section 3.3 are sufficient to show that $\sigma \leq \sigma^{FB}$. In this case, the principal always reduces project level risk to reduce the volatility of the agent's continuation utility.

The structure of the implementation is unchanged, with different values of $\{\phi, \theta\}$, as long as $\{\beta, \Sigma\}$ is invertible as in Proposition 4.

B.2 A General Private Benefits Function

We can also extend the model in Section 2 to consider general, non-linear specifications for the agent's private benefits from resource misallocation. In the main model, the linearity of private benefits implies that the agent derives a constant marginal benefit from resource misallocation. Combined with the decreasing returns to scale of $f(K)$, this constant benefit can be interpreted as "there is a shortage of good projects but no shortage of bad projects". Generalizing to a non-linear private benefits function can be interpreted to mean that bad projects are variable in their utility.

We can repeat the arguments from Section 3 to derive the incentive compatibility condition with a generic private benefit. Let $\Lambda(\hat{K}, K)$ denote the generic private benefit from resource misallocation and be $\mathcal{C}^2$ with the appropriate shape/concavity. Let $\Lambda_{\hat{K}}(K, K)$ denote $\Lambda_{\hat{K}}(\hat{K}, K)|_{\hat{K}=K}$, which is the principal's desired marginal value of private benefits. Then,

$$\beta(\sigma_t, K_t) = \frac{\Lambda_{\hat{K}}(K_t, K_t)}{f'(K_t)} \times \frac{1}{\mu(\sigma_t) - \mu'(\sigma_t)\sigma_t}. \tag{B.3}$$

The optimal contract will differ from the one in Section 3 in two ways. The first is that required incentives may increase or decrease with project size:

$$\frac{\partial}{\partial K}\beta(\sigma, K) = \left(\frac{\Lambda_{\hat{K}K} + \Lambda_{\hat{K}\hat{K}}}{\Lambda_{\hat{K}}} - \frac{f''(K)}{f'(K)} \right) \beta(\sigma, K), \tag{B.4}$$

which may be positive or negative, unlike (12). If the marginal private benefit of misallocating resources declines with project size, and the decline is faster than the decline in the marginal productivity of productive use, then the agent's gain from misallocation is smaller when managing a larger project. Consequently, the incentives required to prevent misallocation can be lower as project size grows. An example uses $\Lambda(\hat{K}, K) = \lambda(K - \hat{K})/K$ over a limited range of K, with $f(K) = K^\alpha, 0 < \alpha < 1$. This implies $\beta_K = -\frac{\alpha}{K}\beta < 0$. One can interpret the assumed Λ to mean that the agent's private benefit depends on the *fraction*, instead of level, of misallocated resources.

Second, the fact that β may decrease in K implies the possibility of over-investment in the optimal contract, or $K(W) > K^{FB}$. The intuition closely follows the intuition for overly risky and overly prudent project choice: as W moves away from W^C, the principal seeks to lower the volatility of W. If β is decreasing in K, the principal may achieve a lower level of continuation volatility with a high K and low β. Consequently, when the principal wants to reduce the volatility of the agent's continuation utility, she may find it optimal to either decrease total cash flow volatility (lower σ and/or lower K) or decrease incentives (higher σ and/or higher K), resulting in different combinations of overly prudent/risky project choices and under/over investment in the equilibrium.

Finally, we have been focusing on optimal contracts following Definition 1, i.e. a contract that eliminates resource misallocation in equilibrium because misallocation is inefficient.

We can consider the case in which misallocation is efficient: for example when $\Lambda(\hat{K}, K) = \lambda(\hat{K} - K)$ with $\lambda > r$. However, this specification changes the right boundary of $F(W)$ in an uninteresting way. If resources used for the agent's private benefit are $\delta \leq \overline{\delta}$, the principal's HJB equation becomes

$$rF(W) = \max_{K,\sigma,\delta} \Big[f(K)\mu(\sigma) - rK - r\delta \ + (\gamma W - \lambda\delta)F'(W) + \frac{1}{2}\beta^2(K,\sigma)f^2(K)\sigma^2 F''(W) \Big].$$

The principal pays private benefits whenever the agent's continuation utility exceeds W^S, defined by $F'(W^S) = -\frac{r}{\lambda}$, instead of the current boundary condition, $F'(W^C) = -1$.

B.3 Dimensionality

Our model uses two-dimensional incentives (β and Σ) to implement a two-dimensional agent's choice (K and σ). What would happen if the agent had more choices, particularly an effort choice? We have assumed that the agent implements a portfolio of projects with a particular risk-return tradeoff, but what if the agent could work to improve the tradeoff? What if the relationship between K and σ were not separable? We might assume that, instead of (1), we have

$$dY_t = \alpha(K_t, \sigma_t, e_t)dt + \varphi(K_t, \sigma_t)dZ_t, \tag{B.5}$$

where e_t is effort. The agent's private benefits function might then be $B(K_t - \hat{K}_t, e_t)$. In this setup, the principal then has two controls (β and $\Sigma = \varphi(K_t, \sigma_t)$) to implement a three-dimensional choice by the agent. The incentive compatibility condition (7) becomes

$$\{K, \sigma, e\} = \arg\max_{\hat{K},\hat{\sigma},\hat{e}} \Big[\beta\alpha(\hat{K}, \hat{\sigma}, \hat{e}) + B(K - \hat{K}, \hat{e})\Big] \tag{B.6}$$

under the constraint that $\Sigma = \varphi(\hat{K}, \hat{\sigma})$.

Because the principal only has two controls, but the agent has three choices, the set of choices the principal can actually implement is a two-dimensional curve in a three-dimensional space. Depending on the functional forms of α and φ, this implementable set might be very far from the first-best. Informally, the principal can condition on output to provide incentives that resources be used efficiently and that total risk is as desired, but her controls are not more granular than that; adding an effort choice need not increase the size of the set of implementable actions.

A key insight of our model of volatility is that it represents a unique setting that cannot be captured by an economically reasonable hidden effort model, such as a simple variation of DeMarzo and Sannikov (2006). One might think that our model could be made simpler by giving the principal control over volatility and resources and also giving the agent a hidden effort choice over drift. For example, using $dY_t = f(K_t)e_t dt + \Sigma_t dZ_t$ instead of (1). However, this change of variables would put all the economics into the private benefits function in a completely uninterpretable way.[32] Instead, we give the agent a second choice over per-unit-of-resources volatility, give the principal a second control over total volatility, and impose economically reasonable assumptions on the production function $f(K)\mu(\sigma)$. These choices are structured so that equilibrium output could be obtained with positive probability under both the agent's true action and under the principal's desired action, thus allowing us to use the martingale methods of Sannikov (2008). In sum, our model demonstrates that agency problems over the composition of volatility are possible and interesting – total cash-flow volatility (Σ), project volatility (σ), and the agent's continuation value volatility ($\beta\sigma$) are all meaningfully different.

[32]In the example, we would have $e_t = \frac{f(\hat{K}_t)}{f(K_t)}\mu\left(\frac{\Sigma_t}{f(\hat{K}_t)}\right)$ as the hidden action instead of $\hat{K}_t$. The agent's private benefits function becomes $B(e, K, \Sigma) = \lambda\left(K - \hat{K}(e, K, \Sigma)\right)$, where $\hat{K}(e, K, \Sigma)$ is an inversion of $e(\hat{K}, K, \Sigma)$. Then, first, $B(e, K, \Sigma)$ has unsigned derivatives (i.e. B_K and B_Σ must change signs at arbitrary point over the relevant range of the problem). Second, even the level of private benefits is difficult to assess because the constraint $B(e(K, K, \Sigma), K, \Sigma) = 0$ is a technological constraint that must be imposed exogenously and has no clear meaning.