capstract:2611.00049v1 [q-fin.GN] 20 Nov 2026

Audibility and Debt Capacity under Contingent Cash Sweeps

Natalie E. Renwick
Department of Finance, Ravelin College of Commerce

November 20, 2026

Abstract

We study a same-day corporate note with a pledge of employee swear-jar receipts. Operating incidents raise the incidence of chargeable speech but move staff onto a restricted bridge where a witness is less likely to hear it. In a six-state model with exact loss enumeration, a witnessed sweep raises debt capacity from 196 to 228 dollars relative to a cash reserve with the same mean contribution under an expected-loss limit. An additional expected-shortfall limit reverses the ranking and lowers capacity to 189 dollars. Pooling four desks on a single incident bridge raises expected shortfall per desk from 87.12 to 109.99 dollars relative to local channels, despite identical mean collections. The pooled structure has fewer missed payments and larger losses in the tail. A standby witness restores pledge value only to the extent that the assignment changes the audibility of speech during closed shifts.

Keywords: corporate liquidity, cash sweeps, debt capacity, expected shortfall, witness access

1 Introduction

An employee who swears on the billing floor owes two dollars to a jar. An employee who swears while muted on the incident bridge owes nothing unless someone on the floor hears it. Both employees may have seen the same failed payment batch. Only one can produce a receivable that the firm can sweep into its debt service account.

A corporate hedge is valuable when it supplies cash the firm would otherwise have to raise at a cost (Froot et al., 1993; Stulz, 1996). A swear jar appears to satisfy this criterion. Frustration rises when receipts are delayed, and collections come from employees rather than from the customers whose payments were delayed. But the jar collects on witnessed speech. As an incident spreads, employees leave the ordinary operating channel and move to a bridge on which microphones are muted by default. The same incident that impairs operating cash can change the conditions under which the supplemental cash is payable.

We study a single-shift note with a standing instruction to sweep employee fines at maturity. We distinguish the occurrence of a chargeable utterance from the conditions under which it becomes a collectible charge. This distinction matters in two comparisons. First, a pledged jar can support more principal than a same-mean cash reserve under an expected-loss limit, and less principal once the lender adds a tail limit. Second, collecting from four desks on a single bridge produces the same expected receipts as collecting from four desks on separate bridges. The distributions of debt recovery differ because a single mute instruction can impair all four jars at once.

This is a structural example with specified inputs. Its antecedents are work on basis risk in contingent coverage (Doherty and Richter, 2002) and the distinction between average exposure

and exposure under stress (Basel Committee on Banking Supervision, 2024). Here the divergence comes from the production of evidence for the charge. A more severe incident can generate more events and fewer provable claims.

2 Contract and model

2.1 The charge and the note

The borrower is a billing operation financed through a note due at the end of the shift. Let S denote the severity of the day's billing incident and let X_s be the operating cash available for this note after prior claims in state s. One can interpret the note as a small increment to desk liquidity. All reported amounts are in dollars per desk per shift unless indicated otherwise.

The swear-jar schedule specifies a fixed fine f. For a charge to be booked, a designated witness must hear an eligible expression live and attribute it to a speaker. A single occurrence can produce at most one charge. The desk has an internal wallet that settles each assessed charge immediately, with sufficient prefunded employee balances. A standing instruction releases the accumulated receipts into a pledged account before the note matures. Surplus cash, including receipts above the amount needed for payment, is released to the borrower.

There are n speech opportunities in a shift. An opportunity is a short interval associated with one worker in which at most one occurrence can be assessed. In state s, an opportunity produces eligible speech with probability p_s. A shift is in open mode with probability a_s. In open mode, a designated witness has access to the ordinary desk channel. In closed mode, incident personnel use a restricted bridge and microphones stay muted except when the moderator grants the floor. A local witness may still hear the speaker. The conditional booking probabilities given an occurrence are q_1 and q_0, with $q_1 > q_0$. Both include attribution. We take open or closed mode to persist across the modeled opportunities, and speech incidence to be independent of mode given severity.

Write B for the indicator of open mode and K for the count of assessed fines. Conditional on (S, B), the opportunities are independent. The cash available from the sweep is $Y = fK$, with

$$K \mid S = s, B = b \;\sim\; \mathrm{Bin}(n, p_s q_b), \qquad \Pr(B = 1 \mid S = s) = a_s. \tag{1}$$

In particular, the binomial calculation conditions on mode. Using a single binomial with the average booking probability preserves the mean and discards the shift-wide variation in witness access.

2.2 The audibility condition

Let $\bar{q}_s = q_0 + (q_1 - q_0)a_s$. The expected cash sweep is $fnp_s\bar{q}_s$. Between two successive severity states with positive incidence and booking probabilities, the sweep falls precisely when

$$\frac{p_{s+1}}{p_s} < \frac{\bar{q}_s}{\bar{q}_{s+1}}. \tag{2}$$

The mean also conceals the change in the count distribution. Setting $a_s = 0$ eliminates the open component even if a simultaneous increase in p_s keeps the mean sweep stable.

For note face D, the lender's same-day loss is

$$L(D) = \max\{D - X_S - Y, 0\}. \tag{3}$$

This definition measures payment at maturity. It leaves subsequent recovery outside the model. We report the probability of a missed payment, expected loss (EL), and the mean loss in the worst 5% of shifts ($\mathbf{ES}_{95}$). To calculate expected shortfall in a discrete distribution, we take only the fraction of the mass at the cutoff needed to reach 5% (Rockafellar and Uryasev, 2002). We define debt capacity as the greatest whole-dollar face that meets an EL limit of 8. A second policy additionally caps $\mathbf{ES}_{95}$ at 100. These are lender risk limits in dollars at a single horizon.

3 Numerical design

A desk has 240 opportunities per shift and a fine of 2. Booking probabilities q_1 and q_0 are 0.80 and 0.02. The former allows for requests to repeat and attribution errors on an open channel. The latter allows for the witness already being in the room when a remote microphone is muted. The severity grid in Table 1 includes three states in which the unassisted desk pays a face of 200 and three in which it falls short. The distinction between outage and recovery is financial. A desk in recovery has exhausted more of the day's receipts while the incident bridge is still closed for most shifts.

Table 1: Scenario inputs per desk. Cash is measured immediately before the sweep.

Severity	Weight (%)	Cash (USD)	Speech p_s (%)	Open a_s (%)
Routine	50	280	2.5	99
Queue	22	245	5.0	99
Backlog	12	215	10.0	98
Degraded	9	150	24.0	85
Outage	5	100	38.0	40
Recovery	2	60	55.0	5

We enumerate the two binomial components at every severity and place mass on the resulting dollar losses using Eq. (3). The inputs are a scenario grid rather than estimates, so the tables report model quantities without sampling intervals. The grid lets the measurement channel change sharply while speech incidence changes progressively. It covers regimes ranging from reliable access to near-complete loss of the ordinary witness.

We compare the stated contract with an unassisted note and two benchmarks. The reserve adds a fixed balance equal to the unconditional expected sweep. The independent sweep draws from the same unconditional distribution as the witnessed sweep, independently of the desk's severity. The latter calculation preserves the jar's skew and zero mass while removing its dependence on operating cash. The reserve comparison holds the mean terminal contribution fixed, with no assertion about the upfront funding cost of a reserve. Both benchmarks use the same seniority and maturity as the witnessed sweep.

4 Results

4.1 Incidence and collectible cash

As the desk moves from degradation to recovery, expected utterances rise from 57.6 to 132 per shift while expected cash receipts fall from 78.68 to 15.58. The expected latent count in recovery

would be sufficient to cover the 140 operating cash gap on a note of 200 at the contractual fine if every utterance paid. With the observed charge rule in the model, 95% of recovery shifts bring in less than 30. For expected receipts alone to cover the recovery gap, the ordinary witness would have to be available on about 65.4% of recovery shifts. The input is 5%.

The conditional distributions in Figure 1 show where the mean comparison loses information. The near-flat segment in the outage curve reaches roughly 60% well below the 100 needed for payment. It represents shifts in the same severity state in which the witness lacks access. A second jar sharing the same witness leaves this probability unchanged.

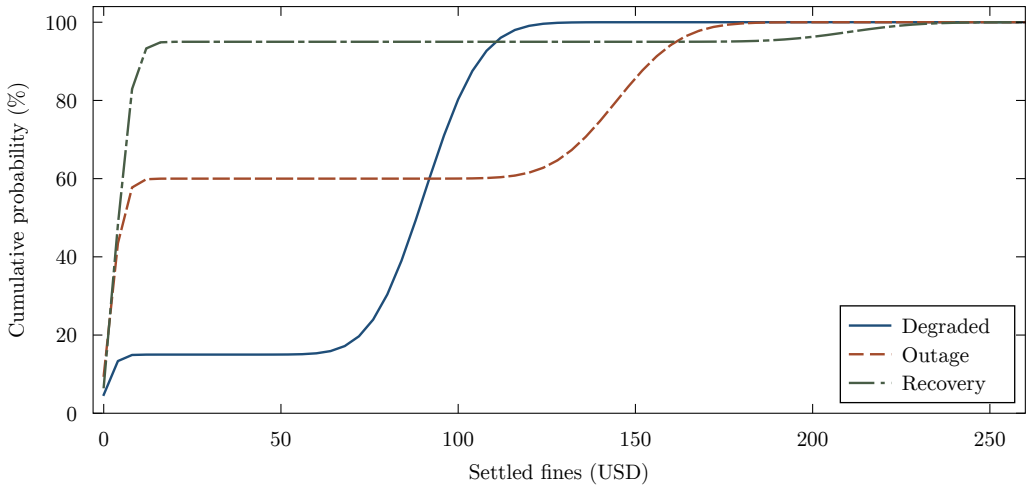


Figure 1: Distribution of collectible receipts in the three states with an operating cash deficit at a face of 200. The curves connect exact cumulative probabilities at four-dollar increments. The cash needed to avoid a missed payment is 50 in degradation, 100 in outage and 140 in recovery.

4.2 Expected loss and the tail limit

Table 2 reports risk on a face of 200 and capacity under each credit policy.

Table 2: Note risk and maximum face per desk under alternative sources of supplemental cash. Risk metrics refer to a face of 200. The last two columns apply the stated lender limits to candidate note faces. Amounts are in USD except for the missed-payment probability.

Supplemental cash	Mean	Miss (%)	EL	ES 95%	EL only	Both limits
None	0.00	16.00	12.30	116.00	173	173
Cash reserve	23.87	16.00	8.48	92.13	196	196
Independent sweep	23.87	14.77	9.03	103.11	193	193
Witnessed sweep	23.87	6.25	6.09	110.00	228	189
Standby, 50% restoration	28.66	3.13	3.05	60.94	259	247
Always open	33.45	0.0023	<0.001	0.002	271	271

The witnessed sweep has the same mean amount as the reserve and cuts expected loss by 2.39 relative to it. It pays most degraded shifts and a large share of outages in full. The missed-payment probability falls from 16.00% to 6.25%. Conditional on a missed payment under the wit-

nessed sweep, average loss is 97.49. Losses are concentrated in closed shifts, and outage and recovery account for 89.4% of expected loss.

Under the EL limit, the lender accepts 32 more dollars of face with the witnessed sweep than with the reserve. Under both limits, it accepts seven fewer. The independent sweep lies between the two benchmarks under the joint limits. Matching both the mean and the unconditional distribution of collections would still give the lender too high a capacity estimate if it severed the relationship between the jar and operating severity. The underwriting decision depends on the probability that staff can be heard on the days when the note needs their money.

5 Pooling across desks

Consider a consolidated note of face mD with access to m equal desks. Severity is common across desks, as when all use the same billing system. We compare a single incident bridge with local channels. In the shared case B is common. In the local case the B_j are conditionally independent and have the same marginal a_s. Utterance opportunities remain independent given severity and the relevant modes. Each desk's expected sweep is identical under the two arrangements.

The ordering extends beyond the scenario grid. Conditional on severity, use one independent uniform draw per opportunity at each desk to construct both of its possible fine counts. Let K_{j0} be its count using the cutoff $p_s q_0$ and let Δ_j be the additional count when the cutoff rises to $p_s q_1$. Each Δ_j is nonnegative. Conditional also on these uniforms, the local total $T_{\text{local}} = \sum_j (K_{j0} + B_j \Delta_j)$ lies between the two endpoints of the shared total $T_{\text{shared}} = \sum_j (K_{j0} + B \Delta_j)$ and has the same mean. The chord inequality for a convex function then gives

$$\mathbb{E}[g(T_{\text{local}}) \mid S = s] \le \mathbb{E}[g(T_{\text{shared}}) \mid S = s]$$

for every convex g. This applies to loss and to each expected excess loss $\mathbb{E}[(L - u)_+]$ for a fixed threshold u. Taking the infimum over u in the expected-shortfall representation preserves the ordering. The probability of missing payment need not follow it.

At four desks, the common bridge gives an expected shortfall of 109.99 per desk, against 87.12 with local channels, even though the probability of any missed payment is lower with the shared bridge (6.25% versus 6.88%). A portfolio report based on arrears frequencies could therefore show an improvement after consolidation. With eight desks the expected-shortfall gap is 26.67 per desk. For a shared bridge, averaging across many desks eliminates idiosyncratic word counts and leaves the two conditional recovery means intact. Figure 2 plots expected shortfall per desk alongside the probability of receiving less than 30 per desk in a recovery shift.

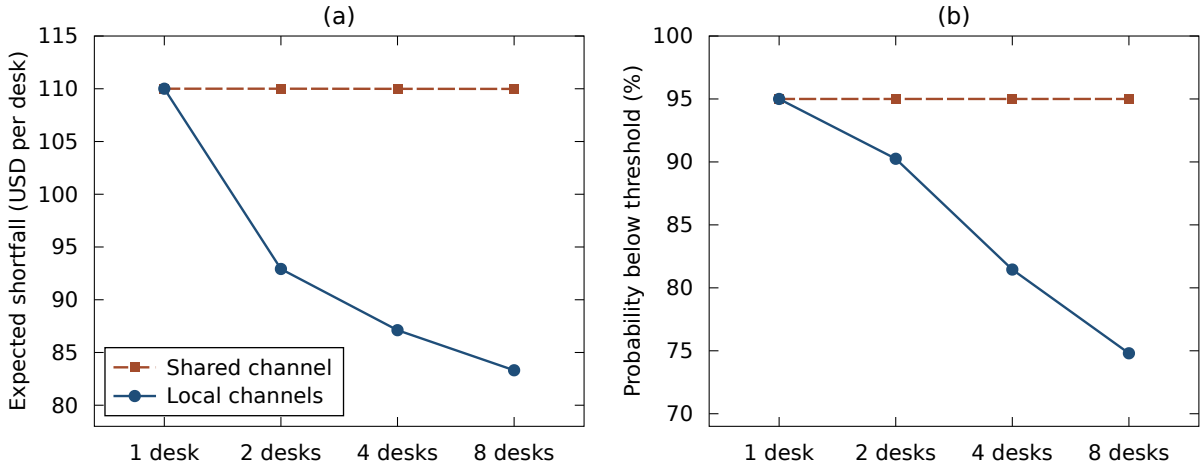


Figure 2: Effect of witness pooling. Left panel reports expected shortfall per desk at 95% for a pooled note of 200 per desk. Right panel reports the probability of an average sweep below 30 per desk, conditional on the recovery state. Expected receipts per desk equal 23.87 in both structures at every portfolio size.

6 A standby witness

A contract can appoint a second witness on a separate channel. Suppose a standby assignment restores open-mode booking with probability r conditional on ordinary closure, without changing speech incidence or operating cash. The open probability becomes $a_s + (1 - a_s)r$. At any fixed face the expected loss is affine in r, since it is a mixture of open- and closed-mode losses. Expected shortfall is piecewise linear and can change slope as different shifts enter its worst 5%. Figure 3 reports both quantities for the note with face 200. Restoring access on 16.23% of closures suffices in the grid to bring its expected shortfall down to 100. At 50% restoration, the drop in expected loss is 3.05 per shift before any witness fee.

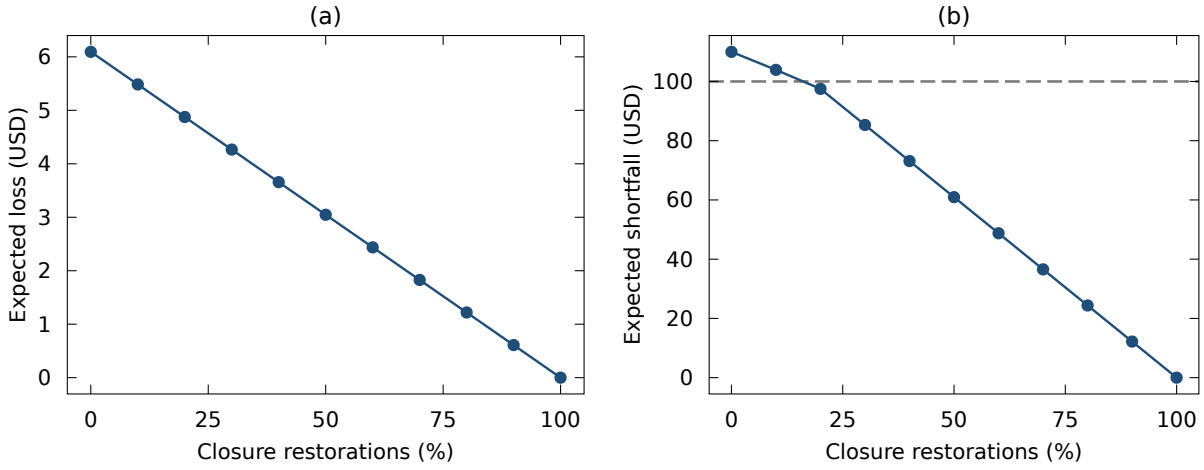


Figure 3: Risk at a face of 200 as a standby witness restores access on a fraction of ordinary closures. The dashed horizontal reference marks the expected-shortfall limit. Separate vertical scales reflect the different magnitudes of the two risk measures.

The value of an intervention depends strongly on when it arrives. For a closed shift, moving from closed to open mode at the start reduces conditional expected loss by 47.70 in degradation, 96.35 in outage, and 134.72 in recovery. It has zero value to the lender in the first three states at this face. Under equal costs per restored shift and severity observed prior to dispatch, a lender minimizing expected loss would call the standby witness to recovery first. The highest-priority assignment is the incident already serious enough that ordinary staff have been told to remain muted.

The calculation treats a standby witness as able to hear a speaker without changing the incident moderator's floor policy and assumes the assignment lasts the shift. Paying someone merely to join the restricted bridge would leave q_0 in effect. A performance test would therefore ask the witness to identify the speaker of a live utterance while the borrower operates in closed mode. A test call performed with the ordinary channel open gives no information about the promised increment in protection.

7 Discussion

The contract assumes immediate collection and sufficient balances in employee wallets. Delayed settlement would introduce a second probability conditional on a valid charge. If wallets run down in severe states, applying the same settlement haircut across all states would give too much credit to cash due during recovery. A borrower also has scope to change speech, monitoring, or ordinary bridge policy once the jar is pledged. The model holds those choices fixed. Induced changes in speech are particularly hard to infer from the ledger, which observes only the product of speech and booking rates.

An empirical implementation would require the incident log, its channel-mode flags, fine timestamps, and an independent count of eligible utterances on a sample of shifts. The administrative ledger alone cannot establish that fewer fines in a severe incident mean fewer utterances. Temporary witness assignment that is independent of incident outcomes within a severity band would allow a comparison of booking rates. Payment timestamps would identify how many documented events generate cash before the note comes due.

The same distinction applies to centralizing the contract. The entities on the legal schedule can have separate payrolls, separate floor managers and separate wallets. A useful register of concentration would also record who can mute their designated witnesses.

Replication

The state grid in Table 1 and the two booking rates fully specify the scenario data. To reproduce the figures, sum the probabilities in Eq. (1) with k as the count index. For the pool with local channels, convolve the one-desk mixture m times within each severity state. For the shared bridge, use mn opportunities conditional on each common mode and mix only after calculating both count distributions. Add operating cash, evaluate Eq. (3) for each support point, and weight by the state probabilities. Compute capacity by increasing the whole-dollar face until a risk limit is breached.

References

Basel Committee on Banking Supervision. Guidelines for Counterparty Credit Risk Management. Technical Report, Bank for International Settlements, 2024. URL https://www.bis.org/bcbs/publ/d588.htm.

Neil A. Doherty and Andreas Richter. Moral Hazard, Basis Risk, and Gap Insurance. *Journal of Risk and Insurance*, 69(1):9–24, 2002. doi: 10.1111/1539-6975.00002.

Kenneth A. Froot, David S. Scharfstein, and Jeremy C. Stein. Risk Management: Coordinating Corporate Investment and Financing Policies. *The Journal of Finance*, 48(5):1629–1658, 1993. doi: 10.1111/j.1540-6261.1993.tb05123.x.

R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional Value-at-Risk for General Loss Distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002. doi: 10.1016/S0378-4266(02)00271-6.

René M. Stulz. Rethinking Risk Management. *Journal of Applied Corporate Finance*, 9(3):8–25, 1996. doi: 10.1111/j.1745-6622.1996.tb00295.x.