Supplier Concentration and Disruption Exposure: A Risk-Scoring Model for Mid-Market Supply Chains

Thomas Anthony Guerriero

Working Paper | 25 November 2026

Abstract

Mid-market firms may buy critical inputs from one or two suppliers because concentration can lower prices and simplify management. The same concentration can turn a supplier's disruption into the buyer's crisis. Large firms can afford specialist tools and staff for supply chain risk assessment; mid-market firms often need something simpler. This conceptual paper proposes a supplier disruption exposure score that a mid-market firm can calculate from information it already holds. The score combines four elements: concentration of spend within each purchased category, measured with the Herfindahl-Hirschman Index as defined in the 2023 United States Merger Guidelines; the criticality of the input to the firm's revenue; the gap between the time needed to recover supply and the protection provided by inventory or qualified alternatives; and credit for warning and recovery capabilities. The model draws on evidence linking announced supply chain problems to prolonged financial underperformance; on propositions linking disruption severity to density, complexity, node criticality, and recovery and warning capabilities; on findings that structural complexity increases disruption frequency; on analysis showing that the best mitigation depends on whether disruptions are frequent and short or rare and long; and on an impact-first exposure model developed with Ford Motor Company. Six propositions address concentration and severity, the trade-off between frequency and severity, impact-first assessment, matching mitigation to disruption profiles, the moderating role of capabilities and score-based sourcing governance. No original empirical data are presented.

Keywords: supply chain risk; supplier concentration; disruption exposure; risk scoring; mid-market firms

Introduction

A mid-market manufacturer may buy from many suppliers, of which a small number supply inputs that the firm cannot easily replace: a custom component, a specialised material or a contract service tied to the firm's own processes. Where purchasing performance is judged mainly on cost, concentrating volume with fewer suppliers is an obvious way to reduce it. Over time, a firm can therefore come to depend heavily on single suppliers for critical inputs without ever making a deliberate decision to accept that risk.

When one of those suppliers is disrupted by a fire, a flood, a cyber incident, insolvency or a failure further up its own supply chain, the effects can be severe. Hendricks and Singhal (2005a) found that firms announcing supply chain glitches experienced large declines in operating performance and did not quickly recover. In a related study, they found average abnormal stock returns of nearly −40% over a period running from one year before to two years after supply chain disruption announcements, together with a significant increase in equity risk (Hendricks & Singhal, 2005b). These studies concern publicly traded firms, but there is no obvious reason to expect smaller firms to be more insulated.

Large firms can address these risks with dedicated supply chain risk teams, specialised software and detailed supplier mapping. Mid-market firms are less likely to have such resources. They need a way of identifying their most dangerous dependencies that is simple enough to be carried out by a purchasing or operations manager using information the firm already holds. This paper proposes such a method: a supplier disruption exposure score.

The paper makes three contributions. First, it brings together research on disruption consequences, disruption severity, structural complexity, mitigation strategy and impact-based exposure assessment into a single scoring logic. Second, it specifies a scoring model, using a standard concentration measure, that mid-market firms can apply without specialised tools. Third, it states six propositions for empirical testing.

The paper is conceptual and reports no original data. The numerical example it includes is hypothetical and is used only to illustrate the calculation. Claims about cited studies are limited to what those studies report.

Literature Review

Why disruptions matter

Hendricks and Singhal (2005a) examined 885 supply chain glitches announced by publicly traded firms and compared the firms' operating performance with that of control firms of similar size and industry. In the year leading up to the announcement, the control-adjusted mean percentage changes in operating income, return on sales and return on assets were −107%,

−114% and −92%, respectively. Relative to controls, firms experiencing glitches reported on average 6.92% lower sales growth, 10.66% higher growth in cost and 13.88% higher growth in inventories. Firms did not quickly recover: in the two years after the announcement, operating income, sales, total costs and inventories did not improve. The authors also found that it did not matter who caused the glitch, what the reason was or what industry the firm belonged to.

In a companion study of 827 disruption announcements made between 1989 and 2000, Hendricks and Singhal (2005b) found average abnormal stock returns of nearly −40% over the period from one year before to two years after the announcement, with much of the underperformance observed in the year before the announcement, on the day of the announcement and in the year after. Equity risk in the year after the announcement was 13.50% higher than in the year before.

These findings indicate that supply chain disruptions are not minor operational inconveniences. They are associated with lasting financial harm, and the harm does not depend on whose fault the disruption was.

Defining and organising supply chain risk

Kleindorfer and Saad (2005) distinguish two broad categories of supply chain risk: risks arising from problems of coordinating supply and demand, and risks arising from disruptions to normal activities, which may stem from natural disasters, strikes and economic disruptions, and the acts of purposeful agents. Their paper provides a conceptual framework reflecting the joint activities of risk assessment and risk mitigation, and draws on data on accidents in the United States chemical industry from 1995 to 2000 to discuss the design of management systems for disruption risk. This paper is concerned with the second category.

Jüttner, Peck and Christopher (2003) sought to clarify the concept of supply chain risk management, provide a working definition and set out a research agenda, comparing the literature with findings from exploratory interviews about practitioners' perceptions of supply chain risk and their risk management strategies. Later reviews show how much the field has grown. Ho and colleagues (2015) reviewed supply chain risk management research published between 2003 and 2013, covering risk definitions, risk types, risk factors and mitigation strategies. Snyder and colleagues (2016) reviewed 180 scholarly works on supply chain disruptions in operations research and management science, organised into six categories: evaluating supply disruptions, strategic decisions, sourcing decisions, contracts and incentives, inventory and facility location.

The breadth of this literature is one reason mid-market firms find it hard to use. The present paper draws selectively on it to construct a tool simple enough for such firms.

What makes a disruption severe

Craighead and colleagues (2007) ask how and why one supply chain disruption is more severe than another. Using a multiple-method, multiple-source research design, they derive six propositions relating disruption severity to three supply chain design characteristics, namely density, complexity and node criticality, and to two mitigation capabilities, recovery and warning. They argue that supply chain disruptions are unavoidable and that all supply chains are therefore inherently risky.

The relevance for supplier concentration is direct. A supplier that provides most or all of a critical input is a critical node in the buyer's supply chain. Craighead and colleagues' framework also identifies two capabilities, recovery and warning, that can reduce the severity of a disruption once it occurs. A scoring model that considers only concentration would miss the difference these capabilities make.

What makes disruptions frequent

Bode and Wagner (2015) address a different question: which supply chain characteristics increase the frequency of supply chain disruptions? Focusing on upstream, supply-side disruptions and drawing on the literature on supply chain complexity, they find that three complexity drivers, horizontal, vertical and spatial complexity, all increase the frequency of disruptions, and that these drivers interact and amplify each other's effects. Their operational definitions of each driver are set out in the original article. In this paper the three drivers are interpreted broadly as the breadth of the supply base, the depth of the supply chain across tiers and its geographic spread; that interpretation is the author's.

Taken together with Craighead and colleagues' (2007) work, this finding implies a tension. Concentrating purchases with fewer suppliers may reduce the structural complexity of the supply base and, on Bode and Wagner's findings, the frequency of disruptions, while increasing the criticality of each remaining supplier and therefore the severity of any disruption that does occur. Broadening the supply base may have the opposite effects. A useful exposure score must capture both sides of this trade-off rather than treating concentration as simply bad.

Choosing a mitigation strategy

Tomlin (2006) studies a single-product setting in which a firm can source from an unreliable supplier and from a reliable but more expensive one. He proves that, in a special case in which the reliable supplier has no volume flexibility and the unreliable supplier has infinite capacity, a risk-neutral firm will pursue a single disruption-management strategy: mitigation by carrying inventory, mitigation by single-sourcing from the reliable supplier, or passive acceptance. A supplier's percentage uptime and the nature of the disruptions, frequent but short or rare but long, are key determinants of the optimal strategy. For a given percentage uptime, sourcing mitigation is increasingly favoured over inventory mitigation as disruptions become less frequent but

longer. A mixed strategy can be optimal if the unreliable supplier has finite capacity or the firm is risk averse.

For mid-market firms, the lesson is that the right protection depends on the profile of the risk. Inventory buffers protect against short interruptions. Qualified alternative sources matter most against long ones.

Assessing impact before probability

Simchi-Levi and colleagues (2015) note that low-probability, high-impact risks are difficult to predict and quantify, so managers may deploy countermeasures suboptimally, leaving some risks exposed while wasting resources on others. In a three-year research engagement with Ford Motor Company, they developed a risk-exposure model that assesses the impact of a disruption originating anywhere in a firm's supply chain. Their approach defers the need to estimate the probability of any specific disruption until after the firm has learned what effect that disruption would have on its operations. Ford applied the model to identify previously unrecognised risk exposures, evaluate pre-disruption mitigation actions and develop post-disruption contingency plans, including where the duration of a disruption is unknown.

This impact-first logic suits mid-market firms well. Such firms may lack the data to estimate the probability that a particular supplier will be disrupted. They can, however, estimate what would happen to their own operations if a supplier stopped delivering.

Measuring concentration

The Herfindahl-Hirschman Index (HHI) is a widely used measure of concentration. The United States Department of Justice and Federal Trade Commission (2023) define it as the sum of the squares of market shares; it is small when there are many small firms, grows larger as concentration increases and reaches 10,000 when a single firm has the whole market. For illustration, the Guidelines note that the HHI for a market of five equal firms is 2,000. The Guidelines treat markets with an HHI greater than 1,800 as highly concentrated for the purposes of merger review.

The present paper borrows the HHI's arithmetic, not its legal thresholds. Applied to a firm's purchases within a category, with each supplier's share of category spend in place of market share, the index summarises how concentrated the firm's sourcing is. A single-sourced category scores 10,000. The antitrust thresholds were not designed for supplier portfolios and are not used as cut-offs here.

Conceptual Framework and Propositions

The supplier disruption exposure score

The proposed score is calculated for each purchased category that supports revenue, and then for the leading supplier within each category. It has four components.

Component 1: Category concentration (C). The HHI of spend shares within the category, divided by 10,000 so that it ranges from near 0 to 1. A category split evenly among four suppliers has a normalised concentration of 0.25; a single-sourced category has 1.0.

Component 2: Criticality (K). The share of the firm's annual revenue that would be lost or deferred if the category's supply stopped, estimated from bills of materials or service dependencies. This reflects Craighead and colleagues' (2007) node criticality and Simchi-Levi and colleagues' (2015) focus on operational impact. A category used in every product scores near 1; a category used in one minor product line scores near 0.

Component 3: Protection gap (G). The ratio of the estimated time to restore supply from the leading supplier or a qualified alternative to the time the firm could continue operating using inventory, pipeline stock and qualified alternative sources, capped at 1. A value of 1 means the firm's protection would run out before, or just as, supply returned; values below 1 mean that existing buffers or alternatives would cover the expected interruption, with a wider margin the lower the value. This component reflects Tomlin's (2006) emphasis on matching mitigation to the duration profile of disruptions.

Component 4: Capability credit (M). A reduction of up to 0.3 for demonstrated warning and recovery capabilities in the relationship with the supplier, such as agreed early notification of problems, visibility of the supplier's own key sub-suppliers and a tested plan to qualify an alternative quickly. This reflects the mitigation capabilities identified by Craighead and colleagues (2007). The cap of 0.3 is a judgement intended to prevent capabilities from masking structural exposure.

The exposure score is:

$$\text{Exposure} = C \times K \times G \times (1 - M)$$

The score ranges from 0 to 1. Multiplying the components means that a category is highly exposed only when it is concentrated, critical and inadequately protected at the same time, which matches the intuition that a single-sourced input of minor importance, or a critical input with ample buffers, is less dangerous than a critical, single-sourced and unbuffered one. The score is intended for ranking and prioritisation within a firm, not as an absolute measure of risk, and its weighting has not been empirically validated.

A hypothetical illustration

Consider a hypothetical mid-market manufacturer reviewing three categories. The figures are invented for illustration only.

- **Machined housings:** one supplier (C = 1.0); used in all products (K = 0.9); four weeks to qualify a new source against two weeks of inventory (G = 1.0, capped); no warning or recovery arrangements (M = 0). Exposure = 1.0 × 0.9 × 1.0 × 1.0 = 0.90.

- **Packaging:** three suppliers with spend shares of 50%, 30% and 20%, giving an HHI of 2,500 + 900 + 400 = 3,800 (C = 0.38); used in all products (K = 0.9); one week to shift volume against three weeks of stock (G = 0.33); no arrangements (M = 0). Exposure = 0.38 × 0.9 × 0.33 × 1.0 ≈ 0.11.

- **Specialty coating:** one supplier (C = 1.0); used in one product line (K = 0.15); twelve weeks to restore against four weeks of stock (G = 1.0); early-warning agreement and mapped sub-suppliers (M = 0.2). Exposure = 1.0 × 0.15 × 1.0 × 0.8 = 0.12.

The housings category is the clear priority. The illustration also shows why concentration alone is a poor guide: the specialty coating is just as concentrated as the housings but far less exposed, because less revenue depends on it.

Proposition 1: Concentration and severity

Craighead and colleagues (2007) relate disruption severity to node criticality and other design characteristics. A supplier that holds most of a firm's spend in a critical category is a critical node. When such a supplier is disrupted, there is less unaffected capacity elsewhere in the supply base to absorb the shortfall.

Proposition 1. For mid-market firms, the concentration of spend within a critical purchased category, measured by the category HHI, will be positively associated with the severity of the operational impact when the leading supplier in that category is disrupted.

Proposition 2: The frequency–severity trade-off

Bode and Wagner (2015) found that complexity drivers increase the frequency of supply-side disruptions. Broadening the supply base may therefore increase the number of disruptions a firm experiences, even as it reduces the severity of each. Concentration has the reverse profile. A firm's total exposure reflects both.

Proposition 2. For mid-market firms, the relationship between the number of suppliers in a critical category and total disruption losses will be curvilinear: losses will be highest under single sourcing because of severity, will fall as a small number of qualified alternatives is added, and will rise again as the supply base becomes broad enough that additional disruption frequency outweighs further reductions in severity.

Proposition 3: Impact-first assessment

Simchi-Levi and colleagues (2015) show that an approach that assesses operational impact before estimating probability helped Ford identify previously unrecognised exposures. Mid-market firms may lack reliable data on supplier disruption probabilities, so probability-first methods may lead them to overlook dependencies that seem unlikely to fail.

Proposition 3. Mid-market firms that assess supplier exposure by first estimating the impact of each supplier's disruption on their own operations will identify more high-exposure dependencies than firms that begin by estimating the likelihood of supplier failure.

Proposition 4: Matching mitigation to the disruption profile

Tomlin's (2006) analysis implies that the longer and rarer the plausible disruption, the stronger the case for an alternative source over a stock buffer. A mid-market firm may rely on safety stock as its main protection because it is familiar and quick to arrange. Where the plausible disruptions are long, such as the loss of a sole supplier's plant, inventory may run out long before supply resumes.

Proposition 4. Among mid-market firms, those that match their mitigation to the expected duration of plausible disruptions, using inventory for short interruptions and qualified alternative sources for long ones, will incur lower disruption losses than those that rely mainly on inventory regardless of disruption profile.

Proposition 5: Capabilities as moderators

Craighead and colleagues (2007) identify recovery and warning as mitigation capabilities related to disruption severity. Early warning from a supplier gives the buyer time to secure stock or alternatives; recovery capability shortens the interruption. These capabilities do not remove the structural exposure created by concentration, but they may reduce its consequences.

Proposition 5. The positive association between category concentration and disruption severity will be weaker for mid-market firms that have established warning arrangements with their leading suppliers and tested plans for recovering supply.

Proposition 6: Score-based sourcing governance

Hendricks and Singhal (2005a, 2005b) found that disruptions were associated with lasting damage to operating performance and share prices. Yet in some mid-market firms, sourcing decisions are made on price and quality, with disruption exposure considered only after a problem arises. Reviewing an exposure score at the point of sourcing decisions would bring that consideration forward.

Proposition 6. Mid-market firms that require a supplier disruption exposure review before awarding or renewing contracts in critical categories will experience fewer severe disruption events over time than firms that consider exposure only after disruptions occur.

Discussion

Concentration is a choice with two sides

The model does not treat supplier concentration as inherently wrong. Concentration can bring lower prices, deeper relationships and simpler management, and Bode and Wagner's (2015) findings suggest that a broader, more complex supply base may experience more frequent disruptions. The aim is to make the cost side of concentration visible so that it can be weighed against the benefits. A firm may reasonably choose to single-source a critical input if it holds sufficient buffers, has a tested recovery plan or receives credible early warning. The score shows when those protections are missing.

Why multiply the components

The multiplicative form reflects the view that exposure arises from the combination of conditions rather than any one of them. An additive score would rate a single-sourced input of trivial importance as moderately risky simply because of its concentration, which would distract attention from genuinely dangerous dependencies. The cost of the multiplicative form is sensitivity: an error in any single component, especially criticality, changes the score proportionately. Firms using the score should therefore take care with the criticality estimate and should review the ranking for results that contradict experienced judgement.

Impact as the starting point

The approach described by Simchi-Levi and colleagues (2015) defers probability estimation until after impact is understood. The proposed score follows this logic by omitting probability altogether. This is deliberate. Mid-market firms may lack defensible probability estimates for supplier disruptions, and including weak estimates could create false precision. A firm that has good information about a particular supplier's vulnerability can use it to adjust priorities once the impact-based ranking is complete.

Borrowing a measure, not a standard

The HHI is used here because it is widely understood, easy to calculate and summarises concentration in a single number. The thresholds used in merger review, however, were designed to assess competition in markets, not the resilience of a buyer's supply base (U.S. Department of Justice & Federal Trade Commission, 2023). There is no reason to assume that an HHI of 1,800 marks a meaningful boundary for supply risk. The model therefore uses the normalised index as a continuous input rather than applying antitrust cut-offs.

Hidden concentration

Spend shares capture concentration among a firm's direct suppliers only. Several direct suppliers may depend on the same sub-supplier, material source or region, so a category that appears diversified may in fact be concentrated further up the chain. Bode and Wagner's (2015) attention

to vertical and spatial complexity and Craighead and colleagues' (2007) attention to supply chain design both point to this limitation. The capability credit partly addresses it by rewarding visibility of key sub-suppliers, but a firm that suspects hidden concentration should investigate directly.

Where the model fits in a wider programme

Ho and colleagues (2015) and Snyder and colleagues (2016) document a wide range of approaches to supply chain risk, including contracts, inventory policy, facility location and strategic decisions. The proposed score is not a substitute for these. It is a triage tool, meant to tell a mid-market firm where to direct its limited attention first. Once the most exposed categories are identified, the firm can draw on the broader repertoire of mitigation approaches to address them.

Practical Implications for Mid-Market Operations and Procurement Leaders

The model suggests the following steps. They are general observations for operational planning, not professional advice for any particular firm.

Calculate concentration from spend data. Many firms can extract spend by supplier and category from their purchasing or accounting systems. Squaring and summing each supplier's percentage share within a category gives the HHI.

Estimate criticality from the products that depend on each input. Bills of materials, service dependencies and revenue by product line allow a rough estimate of how much revenue each category supports.

Compare recovery time with protection time. For each concentrated, critical category, operations staff can estimate how long it would take to restore supply after a serious supplier disruption and how long current stock and qualified alternatives would last.

Match protection to the profile. Where plausible disruptions are short, inventory may be adequate. Where they could last months, qualifying an alternative source, even at modest volume, may be the more effective protection.

Build warning and recovery into supplier relationships. Agreements on early notification, visibility of key sub-suppliers and pre-agreed recovery steps can reduce the severity of a disruption without changing the sourcing structure.

Review exposure when sourcing decisions are made. Adding an exposure check to the approval of new contracts and renewals in critical categories brings risk into the decision before it becomes a problem.

Revisit the ranking regularly. Product mix, suppliers and inventory policies change. A periodic recalculation keeps the ranking current.

Limitations and Future Research

This paper is conceptual and presents no original data. The scoring model, its components and its weights are the author's construction and have not been validated. The hypothetical example illustrates the arithmetic, not real outcomes.

Several cited studies concern large, publicly traded firms (Hendricks & Singhal, 2005a, 2005b) or a single large manufacturer (Simchi-Levi et al., 2015). Their findings may not transfer directly to mid-market firms, whose supplier relationships, bargaining power and resources differ. Tomlin's (2006) results are derived from a stylised model with two suppliers and a single product, and the conditions under which a single strategy is optimal are specific.

The model relies on estimates, especially of criticality and recovery time, that may be uncertain. Its multiplicative form amplifies errors in any one component. It considers direct suppliers and may miss concentration further up the supply chain. It also omits probability by design, which means that a firm could devote attention to a highly exposed dependency that is, in fact, very unlikely to fail.

Future research could test the propositions using data on disruptions experienced by mid-market firms, compare the ranking produced by the score with experienced managers' judgements and with realised losses, and examine whether the curvilinear relationship proposed in Proposition 2 holds across industries. Studies could also calibrate the capability credit and develop simple methods for detecting hidden concentration in sub-tier suppliers.

Conclusion

Supplier concentration can offer mid-market firms real benefits, but it can also leave them exposed to disruptions that research associates with lasting financial harm. This paper has proposed a supplier disruption exposure score that combines category concentration, measured with the Herfindahl-Hirschman Index, with the criticality of each input, the gap between recovery time and existing protection, and credit for warning and recovery capabilities. Drawing on research on disruption consequences, severity, frequency, mitigation strategy and impact-first assessment, the model is designed to be simple enough for firms without specialist risk staff while capturing the trade-off between the frequency and severity of disruptions. Its six propositions provide a basis for testing whether such a score helps mid-market firms direct scarce attention to the dependencies that matter most.

References

Bode, C., & Wagner, S. M. (2015). Structural drivers of upstream supply chain complexity and the frequency of supply chain disruptions. *Journal of Operations Management, 36*(1), 215–228. https://doi.org/10.1016/j.jom.2014.12.004

Craighead, C. W., Blackhurst, J., Rungtusanatham, M. J., & Handfield, R. B. (2007). The severity of supply chain disruptions: Design characteristics and mitigation capabilities. *Decision Sciences, 38*(1), 131–156. https://doi.org/10.1111/j.1540-5915.2007.00151.x

Hendricks, K. B., & Singhal, V. R. (2005a). Association between supply chain glitches and operating performance. *Management Science, 51*(5), 695–711. https://doi.org/10.1287/mnsc.1040.0353

Hendricks, K. B., & Singhal, V. R. (2005b). An empirical analysis of the effect of supply chain disruptions on long-run stock price performance and equity risk of the firm. *Production and Operations Management, 14*(1), 35–52. https://doi.org/10.1111/j.1937-5956.2005.tb00008.x

Ho, W., Zheng, T., Yildiz, H., & Talluri, S. (2015). Supply chain risk management: A literature review. *International Journal of Production Research, 53*(16), 5031–5069. https://doi.org/10.1080/00207543.2015.1030467

Jüttner, U., Peck, H., & Christopher, M. (2003). Supply chain risk management: Outlining an agenda for future research. *International Journal of Logistics Research and Applications, 6*(4), 197–210. https://doi.org/10.1080/1367556031000162701 6

Kleindorfer, P. R., & Saad, G. H. (2005). Managing disruption risks in supply chains. *Production and Operations Management, 14*(1), 53–68. https://doi.org/10.1111/j.1937-5956.2005.tb00009.x

Simchi-Levi, D., Schmidt, W., Wei, Y., Zhang, P. Y., Combs, K., Ge, Y., Gusikhin, O., Sanders, M., & Zhang, D. (2015). Identifying risks and mitigating disruptions in the automotive supply chain. *Interfaces, 45*(5), 375–390. https://doi.org/10.1287/inte.2015.0804

Snyder, L. V., Atan, Z., Peng, P., Rong, Y., Schmitt, A. J., & Sinsoysal, B. (2016). OR/MS models for supply chain disruptions: A review. *IIE Transactions, 48*(2), 89–109. https://doi.org/10.1080/0740817X.2015.1067735

Tomlin, B. (2006). On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science, 52*(5), 639–657. https://doi.org/10.1287/mnsc.1060.0515

U.S. Department of Justice & Federal Trade Commission. (2023). *Merger guidelines* (issued December 18, 2023). https://www.justice.gov/d9/2023-12/2023%20Merger%20Guidelines.pdf