THE ARCHITECTURE AND GOVERNANCE OF COMPLEX ARTIFACTS

IE UNIVERSIDAD

TESIS DOCTORAL/ DOCTORAL DISSERTATION

La arquitectura y el gobierno de los artefactos complejos
/ The Architecture and Governance of Complex Artifacts

Akhil Bhardwaj

SEGOVIA, 2019

IE UNIVERSIDAD

TESIS DOCTORAL/ DOCTORAL DISSERTATION

La arquitectura y el gobierno de los artefactos complejos
/ The Architecture and Governance of Complex Artifacts

Akhil Bhardwaj

Doctoral Thesis Advisor: Mikko Ketokivi

ABSTRACT

The world is populated with complex artifacts, i.e., purposefully designed complex systems encompassing both technology and human actors. Identifying the root cause of their failures challenges the limits of managers' comprehension, and impedes efficient contracting. In response, managers adopt a reductionist approach predicated on the assumption that complex technical systems can be effectively decomposed, i.e., partitioned according to the systems architecture. To evaluate the validity of the assumption of such partitioning for these complex technical systems, I consider locomotive engine failures. These locomotive engines are complex technical systems, and I find they are not decomposable. To facilitate more efficient contracting for these complex technical systems, I suggest that contracts be aligned to identified boundaries of decomposability. Because some complex technical systems are not decomposable, managers are sometimes unable to identify definitively the root causes for failures of these systems. When mangers are unable to identify the root causes of failure, there is a high likelihood of disagreement between transacting parties. Managers that default to invoking opportunism to explain ill-understood failures of complex technical systems, and enforce the letter of the contract, are likely to cause contract termination. Because failures of complex technical systems can cause catastrophic damage, the interests of transacting to mitigate failure is posited to be a shared goal. To attenuate the contracting problem that arises due to the inability of mangers to identify the root causes of artifact failures, I develop a framework that addresses managers' contractual choices to coordinate better high-reliability artifacts. I submit that aligning decision rights with expertise, contracts that empower those with greater local knowledge, and unanticipated outcomes in the technical system that

receive subsequent managerial attention can lower the likelihood of artifact failures occurring due to contracting and technical reasons.

RESUMEN

El mundo está poblado por artefactos complejos, es decir, sistemas complejos diseñados a propósito que abarcan tanto a la tecnología como a actores humanos. La identificación de la causa raíz de sus fallas desafía los límites de la comprensión de los gerentes e impide la contratación eficiente. En respuesta, los gerentes adoptan un enfoque reduccionista basado en el supuesto de que los sistemas técnicos complejos se pueden descomponer de manera efectiva, es decir, se pueden dividir de acuerdo con la arquitectura de los sistemas. Para evaluar la validez del supuesto de tal partición para estos sistemas técnicos complejos, considero las fallas de los motores de locomotoras, que son sistemas técnicos complejos. Me parece que los motores de locomotoras no son descomponibles. Para facilitar una contratación más eficiente para estos sistemas técnicos complejos, sugiero que los contratos estén alineados con los límites identificados de descomposición. Debido a que algunos sistemas técnicos complejos no son descomponibles, los gerentes a veces no pueden identificar definitivamente las causas raíz de las fallas de estos sistemas. Cuando los gerentes no pueden identificar las causas de la falla, existe una alta probabilidad de desacuerdo entre las partes que realizan las transacciones. Es probable que los gerentes que omiten invocar el oportunismo para explicar las fallas mal entendidas de los sistemas técnicos complejos causen la terminación del contrato. Debido a que las fallas de los sistemas técnicos complejos pueden causar daños catastróficos, los intereses de las transacciones para mitigar las fallas se consideran un objetivo compartido. Para abordar el problema de contratación que surge debido a la incapacidad de los gerentes para identificar las causas de la falla de los artefactos, desarrollo un marco que aborde las decisiones contractuales de los gerentes para coordinar mejores artefactos de alta confiabilidad. Afirmo que alinear los derechos de decisión con

la experiencia, los contratos que empoderan a aquellos con mayor conocimiento local y los resultados imprevistos en el sistema técnico que reciben atención gerencial subsiguiente pueden reducir la probabilidad de que ocurran fallas de artefactos debido a razones técnicas y de contratación.

ACKNOWLEDGEMENTS

A task as Brobdingnagian as writing this doctoral thesis could not have been accomplished without considerable help, and I would be remiss not to thank those who helped me.

Mikko Ketokivi, for being a brilliant mentor. Your influence has profoundly shaped how I think and approach scholarship, and I consider myself very lucky to have had you as my adviser. Also, thank you for the conversations about Hitch, Harris, and others. And finally, thank you for generally helping me keep my chin up! I suspect without your mentorship, I may have faltered. It’s been an absolute pleasure!!

Joe Mahoney, for hosting me at Urbana-Champaign multiple times, for your guidance, endless patience, conversations about research, scholarship and general advice. Thank you for being open minded, and being willing to be convinced. You set a standard for me, one that had me returning to write, rewrite, and rethink. I am a better student as a result.

To Dan Belisle, you taught me almost everything I know about railroads, and much more. I’ve enjoyed every conversation we’ve had.

Thank you to my committee members: Matthias Siefert for your support and optimism, it meant a lot. Thank you Saku Mantere and Kiron Ravindran for your feedback and advice. My dissertation is better as a result of your input.

To my fellow travelers: Shijith, Claudio, Mohammed, Sreeya, and others, for the indelible memories and friendship. I summon almost all the Spanish I learned: me casa es su casa!!

To the folks at IE: Laura, Maria, and others, thank you. I presume the printer in the office will work smoothly in my absence!!

To my friends in India: Doc, Arjun, Raiya, Sunny, Shweta, and others. Each of you helped me in your own way, thank you.

To my parents, Suresh and Minakshi Bhardwaj, my thesis literally would not have been written without your support. And thank you for your endless patience and love.

And last but not the least, thank you to my wife Anastasia. For sharing my tribulations and joys, for your untrammeled support, and honesty, and love.

LIST OF CONTENTS

LIST OF TABLES AND FIGURES

Tables

Chapter 1

Chapter 2

Figures

Chapter 1

Chapter 2

INTRODUCTION

We inhabit a world where bacteria can cause train derailments, and errant fasteners used for assembly in a turbosuperchargers on locomotive engines can cause massive failures of complex artifacts, i.e., man-made systems designed to "meet human goals and purposes" (Simon, 1996: 3). The underpinning problem confronting managers is – *how to design, operate, and govern complex artifacts for which some form of failure is unpredictable, inevitable, and in some cases, has the potential to cause catastrophic damage?*

Complex artifacts behave in unpredictable and unfathomable ways (Augier & March, 2008), and often the root causes of their failures cannot be traced (Dekker, Cillier, & Hofmeyr, 2011). Yet, some form of root cause analysis persists, and managers assigning (financial) responsibility (MacDuffie, 1997) often use the results of these analyses as *de facto* premises. These root cause analysis techniques assume that the underpinning technical system can be effectively partitioned such that root causes can be traced, yet this assumption about partitioning complex technical systems remains largely untested. Not surprisingly, investigators have found no evidence that the root cause analysis is effective (Wu, Lipshutz, & Pronovost, 2008).

Managers' inability to identify definitively a root cause responsible for the failures of complex artifacts leads to significant uncertainty. Typically, under circumstances characterized by uncertainty stemming from a measurement problem, vertical integration is prescribed as an adaptive response (Williamson, 1985), especially when opportunism may be in play (Klein, Crawford, & Alchian, 1978). However, vertical integration is not always feasible. For example, one of the contacting parties may be a state-owned entity and not for sale, ruling out the possibility of vertical integration. Furthermore, planning within the vertically integrated firm may challenge managers' cognitive limits (Hayek, 1945). Most critically, the vertical integration response is prescribed for uncertainty originating in *human* behavior, and is not designed for adaptive needs that originate elsewhere.

For example, uncertainty arising from the incompletely understood underpinning complex technical system may be a substantial source of friction between transacting parties leading to honest disagreements. These honest disagreements may be the result of managers being confronted with ill-structured problems originating in the underpinning complex technical system during the management of complex artifacts, a possibility that has been overlooked by theories of contracting such as Transaction Cost Economics, which tend to focus on human behavioral uncertainty (Alchian & Woodward, 1988; Crook, Combs, Ketchen, & Aguinis, 2018). Transaction Cost Economics implicitly posits that the complex artifact and underpinning complex technical system is appropriately partitioned such that decision rights are aligned with managers' expertise, and the primary challenge for contracting is the design of incentives to elicit proper manager behavior. Transaction Cost Economics presupposes that the goals of managers from differing organizations are typically misaligned, yet this supposition may not hold in some important cases.

In the context of high-reliability artifacts, managers from differing organizations are very likely to have a joint goal of preventing complex artifact failure. Preventing the failure of high reliability complex artifacts has been the focus of research in safety science (Dekker *et al*. 2011; Weick & Sutcliffe, 2001). However, the extant literature on safety science neglects the contractual aspect of managing high-reliability complex artifacts. Similarly, theories of contracting overlook uncertainty stemming from the underpinning complex technical system. Hence, there remains a gap in our understanding of failures (McMillan & Overall, 2017), and how to mitigate the failures of complex artifacts in general.

In my doctoral dissertation, I examine the architecture and governance of complex artifacts with a goal of mitigating their (technical and governance) failure. I empirically investigate the behavior of (architecturally) nearly decomposable complex technical systems underpinning complex artifacts. Near decomposability is an assumed architectural and behavioral feature of complex systems (Simon, 1962), and acts an implicit premise in theories of contracting (Arrow, 1985; Williamson, 1985). Because near decomposability is incorporated as a premise in current theorizing about complex systems and their failures, I submit that the literature presupposes a far greater understanding of complex system (failures) than managerial understanding of complex technical systems warrant. Consequently, theories of contracting employing the premise of near decomposability focus on uncertainty arising from human behavior while neglecting anomalies and failures originating in the underpinning technical systems. These anomalies and failures are a mainstay feature of complex technical systems (Waldrop, 1993). Due to this tendency of neglecting anomalies and failures of complex technical systems, and assuming these systems are nearly decomposable, theories of design and contracting encourage managers to reason from an empirically unestablished premise of decomposability. When managers reason from the premise that complex systems are ‘nearly’ decomposable, they are exposed to (some) avoidable but unanticipated hazards. As Mill (1882: 981) noted, the introduction of false premises can lead us astray. Therefore, it is critical to investigate empirically if nearly decomposable systems are, in practice, nearly decomposable.

In the first chapter of my dissertation, I evaluate the assumption that nearly decomposable systems are, in practice, nearly decomposable, using India's Right to Information Act of 2005, I obtained detailed failure data on the operation of a fleet of 919 diesel electric locomotives owned and operated by Indian Railways. Between 2007 and 2015, there were 172 failures of diesel engines; failure means that something happened in the engine that led to an unplanned shutdown and brought the locomotive to an unscheduled halt. Based on engineering drawings, I partitioned the diesel engines of the locomotives into five sub-assemblies. In some cases, failure in only one part of the diesel engine was enough to yield a system failure; a failed piston is a good example. However, in the vast majority of diesel engine failures, it was not just one part (or one sub-assembly), but rather two, three, or even five different parts (or sub-assemblies) that failed simultaneously. The analysis of the diesel engine failures revealed that even though these engines are fully decomposable by design (can be completely disassembled if needed), failures often propagate from one part of the diesel engine to another. This propagation of failures in the diesel engine means that "nearly decomposable" systems are not decomposable, i.e., it does not follow that because a complex system is architecturally decomposable it is decomposable in operation.

Using Cox regression, I was also able to ascertain roughly the effect of type of coupling (tight or loose) on time to failure as well as how failure propagates through the linkages of a system, in this case the diesel engine. Predicated on my findings regarding diesel engine failures, I submit that these sorts of complex systems exhibit *failure complexity,* a condition where the interactions between subsystems preclude the ability of managers to establish root cause of failure. The inability to definitively identify the root cause of failure has direct bearing on 'make or buy' decisions, design of contracts, and organizational boundaries – in the absence of such information, contracting is a challenge. To remedy these contracting problems, I suggest aligning purchase decision with subsystem interfaces, which exhibit low failure complexity.

Having determined that locomotive diesel engines exhibit high failure complexity, in the second chapter of my dissertation, I examine whether high failure complexity might lead to disagreements between Indian Railways and its vendors. I posit honest disagreements between Indian Railways and its vendors might occur because the managers' understanding of the underpinning complex technical system is limited, and anomalies (failures) originating in the complex technical system might be misconstrued by managers as signs that the transacting partner is opportunistic. Opportunism is a central behavioral assumption for theories of contracting (Arrow, 1985; Williamson, 1985), and is considered to be a necessary condition for contractual disagreements and contract terminations (Williamson, 1985). I confront opportunism with the empirical case of Indian Railways buyer-supplier contracts. While opportunism is certainly a sufficient condition for contractual disagreements and contract termination, I posit that it is not a necessary one for these undesirable outcomes. My analysis indicates when managers are dealing with ill-structured problems (Simon, 1973), beset with uncertainty originating in the underpinning complex technical system, limits of bounded rationality of the managers is challenged, and the likelihood of disagreements between transacting parties is high. The likelihood of contract termination is high when transacting parties agree about the nature of the world, but the managers' premises about the underpinning complex technical system is empirically flawed, and they enforce the 'letter of the contract.' Specifically, when managers of both transacting parties assume that locomotive engines are fully decomposable, a root cause for engine failure cannot be definitively established, and a transacting party defaults to assigning the cause of failure to opportunism despite the absence of evidence of such behavior, and thereby may terminate the contract. I suggest that *opportunism is not a necessary condition for contract termination*. Predicated on my analysis, I recommend to managers to evaluate the possible occurrence of opportunism on a case-by-case basis.

Having developed a better understanding of some of the problems confronting managers in the management of complex artifacts, the challenge is to develop an approach to contract for high-reliability complex artifacts under Knightian uncertainty (Knight, 1921). An important consideration for mitigating contracting problems associated with high-reliability complex artifact failures is that some technical failures are inevitable and unpredictable, and these artifacts are open systems (Thompson, 1967). That is, the external environment influences their behavior in unpredictable ways. Transaction Cost Economics (Williamson, 1985), which is one of the most prevalent theories of contracts, I submit is not sufficient because it neglects to recognize complexity in its entirety, and overlooks uncertainty arising from the underpinning technical system and socio-technical interactions that occur at the interface of human agents and technology. In addition, Transaction Cost Economics does not consider the focal problem of the prevention of technical system failures. Yet for complex artifacts, like railroads, in addition to mitigating contracting failures, mitigating technical failures is also vital. Aside from economic costs, failures of high-reliability artifacts can lead to loss of life.

In the third chapter of my dissertation, I develop framework to attenuate artifact failure when contracting for high-reliability artifacts in conditions of (Knightian) uncertainty. The challenge of mitigating artifact failures is compounded by the fact that managers are severely limited in their cognitive abilities. On the one hand, these managers have *bounded hindsight*, which limits their ability to discern causal connections between events leading up to artifact failure. On the other hand, these managers have *bounded foresight*, which limits their ability to devise safeguards adequately to mitigate artifact failure originating in non-behavioral factors. Managers are also boundedly rational, which increases their likelihood of being unaware of anomalies due to limits on attention. Managers also have a proclivity to fit those anomalies that they do become aware of into their pre-existing mental schemas (Starbuck & Farjoun, 2005).

The developed framework proposes that when unpredictable artifact failure occurs, managers of transacting parties search for non-behavioral causes of failure rather than defaulting to assigning the cause to opportunism. Managers can design governance structures that promote the joint pursuit of safety by transacting parties, while directing the managers' attention to anomalies and failures regardless of the immediate economic costs, and search for (multiple) explanations for these events. The framework suggests that managers match (inter-disciplinary) expertise with responsibility and decision rights for contract design, and encourages managers to design simpler contracts. This dissertation study offers a technical system-rationale for artifact failure without relying on the premise of behavioral uncertainty (e.g., opportunism).

To summarize, in my dissertation, using the empirical case of locomotive engine failures, I evaluate the theoretical premise of near decomposability, which is often employed in theories of manufacturing, operations and contracting. I find that 'nearly decomposable' systems are not nearly (enough) decomposable, i.e., failures propagate from one subsystem to another, and some systems exhibit high failure complexity, which means managers cannot pinpoint the root cause of failure. I suggest governance structures be aligned with boundaries of near decomposability (of failures). Next, I seek to ascertain if transacting parties may have honest disagreements when contracting for complex artifacts, and identify antecedents for those disagreements. I find that the likelihood of disagreements is high when managers are confronted with ill-structured problems that have their origins in the underpinning complex technical system. In addition, these disagreements may escalate and lead to contract termination when a transacting party attributes opportunism to its transacting partner without evidence of opportunistic behavior by the partner. Finally, I develop a framework for contracting for high-reliability artifacts under conditions of Knightian uncertainty, which is most applicable for high-reliability systems where the fallout from failure may be catastrophic, and the interests of contracting parties is aligned in that they want to mitigate artifact failure. The framework focuses on making the high-reliability artifacts less susceptible to failure, recognizes the severe cognitive limits of managers, and proposes measures such as aligning decision rights with requisite manager expertise and the use of simpler contracts.

INTRODUCCIÓN

Vivimos en un mundo donde las bacterias pueden causar descarrilamientos de trenes, y los sujetadores errantes utilizados para el montaje en sobrealimentadores turbo en motores de locomotoras pueden causar fallas masivas de artefactos complejos, es decir, sistemas hechos por el hombre diseñados para "cumplir con los objetivos y propósitos humanos" 1996: 3). El problema subyacente que enfrentan los gerentes es: *¿cómo diseñar, operar y gobernar artefactos complejos para los cuales algún tipo de falla es impredecible, inevitable y, en algunos casos, tiene el potencial de causar un daño catastrófico?*

Los artefactos complejos se comportan de manera impredecible e insondable (Augier y marzo, 2008), y con frecuencia las causas principales de sus fallas no se pueden rastrear (Dekker, Cillier y Hofmeyr, 2011). Sin embargo, persiste alguna forma de análisis de causa raíz, y los gerentes que asignan responsabilidades (financieras) (MacDuffie, 1997) a menudo usan los resultados de estos análisis como premisas de facto. Estas técnicas de análisis de causa raíz asumen que el sistema técnico subyacente puede particionarse de manera efectiva de tal manera que se puedan rastrear las causas raíz, sin embargo, esta suposición acerca de la partición de sistemas técnicos complejos aún no se ha probado. No es sorprendente que los investigadores no hayan encontrado pruebas de que el análisis de la causa raíz sea efectivo (Wu, Lipshutz y Pronovost, 2008).

La incapacidad de los gerentes para identificar definitivamente una causa raíz responsable de las fallas de artefactos complejos conduce a una incertidumbre significativa. Normalmente, bajo circunstancias caracterizadas por la incertidumbre derivada de un problema de medición, la integración vertical se prescribe como una respuesta adaptativa (Williamson, 1985), especialmente cuando el oportunismo puede estar en juego (Klein, Crawford y Alchian, 1978). Sin embargo, la integración vertical no siempre es factible. Por ejemplo, una de las partes contratantes puede ser una entidad estatal y no para la venta, descartando la posibilidad de integración vertical. Además, la planificación dentro de la empresa integrada verticalmente puede desafiar los límites cognitivos de los gerentes (Hayek, 1945). De manera más crítica, la respuesta de integración vertical se prescribe para la incertidumbre que se origina en el comportamiento humano, y no está diseñada para necesidades adaptativas que se originan en otros lugares.

Por ejemplo, la incertidumbre que surge de un sistema técnico complejo que no se entiende completamente, puede ser una fuente importante de fricción entre las partes en proceso de transacción que conduce a desacuerdos honestos. Estos desacuerdos honestos pueden ser el resultado de la confrontación de los gerentes con problemas mal estructurados originados en el complejo sistema técnico subyacente durante la gestión de artefactos complejos, una posibilidad que ha sido pasada por alto por las teorías de contratación como Economía de costo de transacción, que tienden a centrarse sobre la incertidumbre del comportamiento humano (Alchian y Woodward, 1988; Crook, Combs, Ketchen y Aguinis, 2018). La economía de costo de transacción implícitamente postula que el complejo artefacto y el sistema técnico complejo que lo sustenta están particionados de manera apropiada, de modo que los derechos de decisión estén alineados con la experiencia de los gerentes, y el principal desafío para la contratación es el diseño de incentivos para obtener el comportamiento adecuado de los gerentes. La economía de costo de transacción presupone que los objetivos de los gerentes de diferentes organizaciones generalmente están desalineados, aunque esta suposición puede no ser válida en algunos casos importantes.

En el contexto de los artefactos de alta confiabilidad, es muy probable que los gerentes de diferentes organizaciones tengan un objetivo conjunto de prevenir fallas complejas de artefactos. Prevenir la falla de artefactos complejos de alta confiabilidad ha sido el foco de la investigación en ciencia de la seguridad (Dekker et al. 2011; Weick & Sutcliffe, 2001). Sin embargo, la literatura existente sobre ciencia de la seguridad descuida el aspecto contractual de la gestión de artefactos complejos de alta confiabilidad. Del mismo modo, las teorías de la contratación pasan por alto la incertidumbre derivada del complejo sistema técnico que lo sustenta. Por lo tanto, sigue habiendo una brecha en nuestra comprensión de las fallas (McMillan & Overall, 2017), y cómo mitigar las fallas de los artefactos complejos en general.

En mi tesis doctoral, examino la arquitectura y el gobierno de artefactos complejos con el objetivo de mitigar su falla (técnica y de gobierno). Investigo empíricamente el comportamiento de sistemas técnicos complejos casi descomponibles (arquitectónicamente) que sustentan artefactos complejos. La descomposición cercana es una característica arquitectónica y de comportamiento asumida de los sistemas complejos (Simon, 1962) y actúa como una premisa implícita en las teorías de contratación (Arrow, 1985; Williamson, 1985). Debido a que la descomposición cercana se incorpora como una premisa en la teorización actual sobre sistemas complejos y sus fallas, sostengo que la literatura presupone una comprensión mucho mayor de sistemas complejos (fallas) que la comprensión administrativa de sistemas técnicos complejos. En consecuencia, las teorías de contratación que emplean la premisa de la descomposición cercana se centran en la incertidumbre que surge del comportamiento humano, mientras que se descuidan las anomalías y fallas que se originan en los sistemas técnicos subyacentes. Estas anomalías y fallas son una característica fundamental de los sistemas técnicos complejos (Waldrop, 1993). Debido a esta tendencia a descuidar las anomalías y fallas de los sistemas técnicos complejos, y al suponer que estos sistemas son casi descomponibles, las teorías de diseño y contratación alientan a los gerentes a razonar desde una premisa de descomponibilidad empíricamente no establecida. Cuando los gerentes razonan a partir de la premisa de que los sistemas complejos son "casi" descomponibles, están expuestos a (algunos) peligros evitables pero no anticipados. Como Mill (1882: 981) señaló, la introducción de premisas falsas puede llevarnos por mal camino. Por lo tanto, es crítico investigar empíricamente si los sistemas casi descomponibles son, en la práctica, casi descomponibles.

En el primer capítulo de mi disertación, evalúo la suposición de que los sistemas casi descomponibles son, en la práctica, casi descomponibles, utilizando la Ley de Derecho a la Información de la India de 2005, obtuve datos detallados de fallas en el funcionamiento de una flota de 919 locomotoras eléctricas diésel. y operado por ferrocarriles indios. Entre 2007 y 2015, hubo 172 fallas en los motores diésel; falla significa que algo sucedió en el motor que llevó a un apagado imprevisto y detuvo la locomotora. Basándome en los planos de ingeniería, dividí los motores diésel de las locomotoras en cinco subconjuntos. En algunos casos, una falla en solo una parte del motor diésel fue suficiente para producir una falla en el sistema; Un pistón averiado es un buen ejemplo. Sin embargo, en la gran mayoría de los fallos de los motores diésel, no fue solo una parte (o un subconjunto), sino más bien dos, tres o incluso cinco partes diferentes (o subconjuntos) que fallaron simultáneamente. El análisis de los fallos de los motores diésel reveló que, si bien estos motores son totalmente descomponibles por diseño (se pueden desmontar por completo si es necesario), los fallos a menudo se propagan de una parte del motor diésel a otro. Esta propagación de fallas en el motor diésel significa que los sistemas "casi descomponibles" no son descomponibles, es decir, no se sigue que debido a que un sistema complejo es arquitectónicamente descomponible se pueda descomponer en operación. Al usar la regresión de Cox, también pude determinar aproximadamente el efecto del tipo de acoplamiento (apretado o suelto) en el tiempo de falla, así como la manera en que la falla se propaga a través de los enlaces de un sistema, en este caso el motor diésel. Predicado en mis hallazgos con respecto a las fallas de los motores diésel, sostengo que este tipo de sistemas complejos presentan una complejidad de falla, una condición en la que las interacciones entre los subsistemas impiden la capacidad de los administradores para establecer la causa raíz de la falla. La incapacidad de identificar definitivamente la causa raíz del fracaso tiene una relación directa con las decisiones de "hacer o comprar", el diseño de los contratos y los límites de la organización; en ausencia de dicha información, la contratación es un desafío. Para remediar estos problemas de contratación, sugiero alinear la decisión de compra con las interfaces del subsistema, que presentan una baja complejidad de fallos.

Después de haber determinado que los motores diésel de locomotoras presentan una alta complejidad de falla, en el segundo capítulo de mi disertación, examino si la complejidad de una falla alta podría llevar a desacuerdos entre Ferrocarriles de la India y sus proveedores. Pienso que los desacuerdos honestos entre Ferrocarriles de la India y sus proveedores podrían ocurrir porque los gerentes entienden que el sistema técnico complejo subyacente es limitado, y las anomalías (fallas) que se originan en el sistema técnico complejo pueden ser malinterpretadas por los administradores como señales de que el socio que realiza la transacción es oportunista. El oportunismo es un supuesto de comportamiento central para las teorías de la contratación (Arrow, 1985; Williamson, 1985), y se considera una condición necesaria para los desacuerdos contractuales y la rescisión de contratos (Williamson, 1985). Enfrento el oportunismo con el caso empírico de los contratos entre compradores y proveedores de Ferrocarriles de la India. Si bien el oportunismo es ciertamente una condición suficiente para los desacuerdos contractuales y la rescisión del contrato, creo que no es necesario para estos resultados indeseables. Mi análisis indica cuándo los gerentes se enfrentan a problemas mal estructurados (Simon, 1973), acosados por la incertidumbre que se origina en el complejo sistema técnico subyacente, se cuestionan los límites de la racionalidad limitada de los gerentes y existe una alta probabilidad de desacuerdos entre las partes que realizan transacciones. La probabilidad de rescisión del contrato es alta cuando las partes en la transacción están de acuerdo con la naturaleza del mundo, pero las premisas de los gerentes sobre el complejo sistema técnico subyacente están empíricamente viciadas. Específicamente, cuando los gerentes de ambas partes que realizan transacciones asumen que los motores de locomotoras son totalmente descompuestos, no se puede establecer definitivamente una causa raíz para la falla del motor, y una parte que realiza la transacción por defecto asigna la causa de la falla al oportunismo a pesar de la ausencia de evidencia de tal comportamiento, por lo tanto puede rescindir el contrato. Sugiero que el oportunismo no es una condición necesaria para la rescisión del contrato. Basado en mi análisis, recomiendo a los gerentes que evalúen la posible ocurrencia de oportunismo caso por caso.

Habiendo desarrollado una mejor comprensión de algunos de los problemas que enfrentan los gerentes en el manejo de artefactos complejos, el desafío es desarrollar un enfoque para contratar artefactos complejos de alta confiabilidad bajo la incertidumbre de Knight (Knight, 1921). Una consideración importante para mitigar los problemas de contratación asociados con fallas de artefactos complejos de alta confiabilidad es que algunas fallas técnicas son inevitables e impredecibles, y estos artefactos son sistemas abiertos (Thompson, 1967). Es decir, el ambiente externo influye en su comportamiento de manera impredecible. La economía de costos de transacción (Williamson, 1985), que es una de las teorías más prevalentes de los contratos, sostengo que no es suficiente porque no reconoce la complejidad en su totalidad y pasa por alto la incertidumbre derivada del sistema técnico subyacente y las interacciones socio-técnicas que Ocurren en la interfaz de los agentes humanos y la tecnología. Además, la Economía de costo de transacción no considera el problema central de la prevención de fallas técnicas del sistema. Sin embargo, para los artefactos complejos, como los ferrocarriles, además de mitigar las fallas de contratación, mitigar las fallas técnicas también es vital. Aparte de los costos económicos, las fallas de artefactos de alta confiabilidad pueden llevar a la pérdida de vidas.

En el tercer capítulo de mi disertación, desarrollo un marco para atenuar la falla de artefactos al contratar artefactos de alta confiabilidad en condiciones de incertidumbre (Knightian). El desafío de mitigar las fallas de los artefactos se ve agravado por el hecho de que los gerentes están muy limitados en sus capacidades cognitivas. Por un lado, estos gerentes han limitado la retrospectiva, lo que limita su capacidad para discernir las conexiones causales entre los eventos que conducen a la falla de artefactos. Por otro lado, estos gerentes han limitado la previsión, lo que limita su capacidad para diseñar salvaguardas adecuadas para mitigar la falla de artefactos originada en factores no conductuales. Los gerentes también son limitadamente racionales, lo que aumenta su probabilidad de desconocer anomalías debido a los límites de atención. Los gerentes también tienen una tendencia a adaptarse a esas anomalías de las que se dan cuenta en sus esquemas mentales preexistentes (Starbuck y Farjoun, 2005).

El marco desarrollado propone que cuando se produce una falla impredecible de artefactos, los administradores de las partes que realizan transacciones buscan causas de fallas no conductuales en lugar de dejar de asignar la causa al oportunismo. Los gerentes pueden diseñar estructuras de gobierno que promueven la búsqueda conjunta de la seguridad por parte de las partes que realizan transacciones, mientras dirigen la atención de los gerentes a las anomalías y fallas, independientemente de los costos económicos inmediatos, y buscan (múltiples) explicaciones para estos eventos. El marco sugiere que los administradores combinan la experiencia (interdisciplinaria) con los derechos de responsabilidad y decisión para el diseño del contrato, y alienta a los administradores a diseñar contratos más simples. Este estudio de disertación ofrece una justificación técnica del sistema para la falla de artefactos sin depender de la premisa de la incertidumbre de comportamiento (por ejemplo, oportunismo).

Para resumir, en mi disertación, utilizando el caso empírico de las fallas de los motores de locomotoras, evalúo la premisa teórica de casi descomponibilidad, que a menudo se emplea en las teorías de fabricación, operaciones y contratación. Encuentro que los sistemas "casi descomponibles" no son casi descomponibles (es decir, los fallos se propagan de un subsistema a otro) y algunos sistemas presentan una alta complejidad de fallos, lo que significa que los administradores no pueden identificar la causa raíz del fallo. Sugiero que las estructuras de gobierno se alineen con los límites de la descomposición cercana (de las fallas). A continuación, busco determinar si las partes en la transacción pueden tener desacuerdos honestos al contratar artefactos complejos, e identificar antecedentes para esos desacuerdos. Encuentro que la probabilidad de desacuerdos es alta cuando los gerentes se enfrentan a problemas mal estructurados que tienen su origen en el complejo sistema técnico subyacente. Además, estos desacuerdos pueden escalar y llevar a la rescisión del contrato cuando una parte de la transacción atribuye el oportunismo a su pareja de la transacción sin evidencia de comportamiento oportunista por parte de la pareja. Finalmente, desarrollo un marco para la contratación de artefactos de alta confiabilidad en condiciones de incertidumbre de Knight, que es más aplicable para sistemas de alta confiabilidad donde las consecuencias del fracaso pueden ser catastróficas, y los intereses de las partes contratantes están alineados en lo que quieren mitigar el fallo del artefacto. El marco se enfoca en hacer que los artefactos de alta confiabilidad sean menos susceptibles a fallas, reconoce los severos límites cognitivos de los gerentes y propone medidas como alinear los derechos de decisión con la experiencia requerida de los gerentes y el uso de contratos más simples.

Chapter 1

Inevitable Failures and Ambiguous Root Causes: Contract Design and Supply Chain Management in the Railroad Industry

Abstract

How should managers think of buyer-supplier contracts in a context where product design is stable but where failures are frequent and their root causes ambiguous? When product redesign is not an option, how should managers approach the problem? In this study, I analyze a dataset of 172 system failures in Indian Railways' fleet of locomotive engines. My goal is both to understand how these engines fail as well as the implications for buyer-supplier contracting in a context where subsystems are produced by external vendors. Based on my empirical analysis, I find that the way responsibility is assigned in the warranty processes in particular does not seem to align with the observed failure patterns. Indeed, vendors may sometimes be held responsible for someone else's problems. However, I propose that these distortions are remediable by rethinking the buyer-supplier relationships. In light of my findings, I discuss the implications of understanding system-level failures for the reassessment of warranty processes in particular, and supply chain management more generally.

Keywords: Buyer-supplier relationships; system failure; contracting; warranty; survival analysis, complexity.

“Failure analysis doesn’t always lead us to the correct root cause, or even any cause. In fact, the most common failure identified and reported by the majority of railroads is NDF—that’s short for No Defect Found.”

Dan Belisle, Senior Field Service Engineer, Progress Rail (subsidiary of Caterpillar, Inc.)

1.1. Introduction

In 2012, Indian Railways downgraded one of its largest suppliers of locomotive engine parts from the preferred Tier 1 status and awarded the multi-million-dollar contract to another vendor, who was, in turn, upgraded to Tier 1. Indian Railways concluded that not only had the original vendor been supplying engine components that did not meet the conformance quality specifications, but also there was a reason to believe they had done this knowingly, and therefore, engaged in opportunistic behavior. Surely, a large, established vendor would know whether the quality of its components was intact. Therefore, a vendor’s challenge of the warranty claim could plausibly be interpreted as opportunism. However, Indian Railways quickly realized that the components supplied by the new vendor were failing in a similar way and with similar frequency. Perhaps the components had failed despite not having been defective. Indian Railways corrected the mistake and reinstalled the Tier 1 status of the original supplier.

Readily equating failure with defect without closer investigation is simultaneously illogical and understandable. It is illogical because to infer defect from failure is a glaring *non sequitur*. Failures occur for many reasons, quality defect being only one possibility. However, equating the two is understandable in that it expedites decision-making and the warranty process: it may be economically attractive to forgo a systematic root cause analysis and simply claim warranty. Indeed, it is common to find contractual clauses in Indian Railways’ buyer-supplier contracts that stipulate that any component that fails within a pre-specified number of months will be warranted.

In the interest of expediting repair of locomotives, it is further common for buyers to claim routinely warranties in the case of early failures. The shortcoming of enforcing such principles is that while eschewing a root cause analysis expedites decision-making, one simultaneously fails to address the underpinning technical problem.

In addition to being inter-organizationally understandable, equating failure with defect facilitates decision-making *within* organizations. In his analysis of shop-floor problems in auto assembly, MacDuffie noted that attempts to identify the root cause of the problem were limited: "[p]roblem analysis, such as it is, is almost entirely concerned with assigning financial accountability" (1997: 486). Therefore, managerial and technical concerns are not only distinct; they may be in conflict with one another. Engineering would prefer that resources be allocated to solving the technical problems, which would require pinpointing the root cause. But managers may lean toward enforcing accounting principles that enable swift, unambiguous assignment of costs to specific organizational sub-units; to this end, being able to justify an *apparent* cause is sufficient.

It is equally illogical and understandable to ascribe readily opportunism (Williamson, 1985: 183-189). Recurrent component failures do not necessarily imply the supplier is dishonest, but understandably raises the question why the supplier does not know that its components are of sub-par quality. Indian Railways selects its suppliers after due diligence and based on competence. Thus, dishonesty is an understandable conclusion. It can also serve as an organizational stopping rule: it gives an account of the problem and can serve as a *de facto* root cause. Opportunism therefore serves as a *factual premise* (Simon, 1976) in drawing the conclusion—whether this factual premise is factually correct is ancillary, what matters is that it is treated as such in making the decision.

However, the problem gets even more complicated. While intuitively appealing, root cause analysis has many problems, starting with the assumptions it embraces. These problems are well known and documented in medicine (Wu, Lipshutz, & Pronovost, 2008) and the safety sciences (Dekker, Cilliers, & Hofmeyr, 2011; Leveson, 2004; Peerally, Carr, Waring, & Dixon-Woods, 2017). Volumes of books and articles written on the topic offer technical tools to implement it. However, while there are studies that focus on system improvement (e.g., Tucker, 2007), I was hard-pressed to find empirical studies that support that root cause analysis *actually works*. Indeed, as Tucker and Edmondson (2003) noted, even in a critical environment such as medical care, evidence of its effectiveness is elusive. The relevant criterion is better captured by the question "Does root cause analysis actually reduce the recurrence of problems?" (Wu *et al*. 2008: 685). I am not aware of any empirical studies that have directly addressed this question systematically.

To address the question, one must show that an organizational intervention that introduces root cause analysis actually solves problems and reduces failures. Being able to find *a* root cause is not enough, one needs to identify *the* root cause, and subsequently, engage in actions that address the technical problem and reduce its recurrence. Root cause analysis is an effective tool in assigning responsibility and allocating cost, but the ultimate "payoff" of root cause analysis is to solve technical problems, and not merely to expedite decision-making.

The fundamental premise in root cause analyses is that an unambiguous cause not only exists but that it can also be unambiguously identified. In the case of technical systems, this translates to assuming that a "eureka part" (Dekker *et al*. 2011: 940) that initiated the system failure exists. The best-known example of chasing the eureka part was the tracing of the Challenger disaster in 1986 to the failure of two rubber gaskets ("O-rings") that failed to seal the rocket booster properly at low temperatures, ultimately inducing a catastrophic system failure (Starbuck & Milliken, 1988). However, the causal chain of events still eludes us: the failure of the O-rings may well have been the proximate cause, but we are not aware of an investigation that established the root cause in an empirically rigorous manner. What we do know is that there is no evidence that the O-rings were defective, but that they had not been tested for the temperature at which the launch occurred.

The Challenger disaster is analogous with Mr. Belisle's—a railroad engineer with 35 years of field experience—observation that most of the time engineers in the railroad industry are unable to identify the root cause of the failure: a failure analysis reveals no evidence of any part of the system having been defective. Indeed, the causal processes associated with system (as opposed to component) failures in particular can be so complex that determining what exactly happened and who is at fault, often requires months, even years, of failure analysis, followed by legal proceedings. A case in point, concerns the fatal Lac-Megantic derailment on July 6, 2013, an unattended 74-car freight train carrying inflammable crude oil rolled downhill into the town of Lac-Megantic, Canada, killing 47 people. Three railroad employees were charged with criminal negligence, but were four years later acquitted of all charges. In its Railway Investigation Report R13D0054 (July 6, 2013), Transportation Safety Board of Canada concluded that it was unable to identify the root cause, instead, they identified a list of 18 distinct contributing factors. The acquittal was of little consolation to the victims or the *Montreal, Maine, and Atlantic Railway,* which had already filed for bankruptcy in 2013, citing specifically liability as the reason.

In this dissertation chapter, I ask: When technical systems fail—as they unexpectedly and inevitably do—and the root cause is ambiguous, how should contractual relationships with suppliers be managed? Further, in the spirit of *the mirroring hypothesis* (Colfer & Baldwin, 2016; Sanchez & Mahoney, 1996), I ask, How can the organizational and contractual arrangements mirror the technical realities of the system if these technical realities—*dependencies* (Colfer & Baldwin, 2016: 709) in particular—are ambiguous? I examine this question both theoretically and empirically in a sample of 172 locomotive engine failures. Using India's Right to Information Act of 2005, I obtained data on all system failures that occurred in 2007-2015 in one particular fleet of technically identical locomotive engines at Indian Railways. The dataset is particularly useful for analyzing system failures, because in addition to having a sufficiently large sample of failures, all failures involve the same engine model, assembled at the same location from identical components with the same rated life. Further, the use of the engines is not limited to a specific region in India; instead, they are operated on an "open-loop" basis, which means that any given locomotive operates in a geographically broad area. Human errors are further eliminated in that the engines are largely controlled by on-board computers; the actions of the train operators (drivers) are very limited. At the same time, the ways in which these engines failed is diverse, and the time to failure varies as well. This variance provides an opportunity examine failure patterns in more detail. The aim is to understand how locomotive engines, as systems, fail, and what managers can learn from the failures. Ultimately, this study is not about engineering, product design, or failure prevention—the key implications are for *contract design*. However, I seek to strike a balance by addressing the managerial problem by incorporating the technical realities of system failures.

My findings suggest that even in the case of a relatively simple and well-understood system such as a locomotive diesel engine, establishing the root cause of failure is elusive: systems fail in ways that make it difficult to pinpoint specifically which subsystems were ultimately contributing to the failure. In such circumstances, following the normative prescription of the mirroring hypothesis becomes difficult, not only because the technical realities are unknown, but also because redesign of a dominant engine design is generally not feasible in the first place. However, starting at the mirroring hypothesis' prescription that problem-solving resources be placed "where problems are most likely to appear" (Colfer & Baldwin, 2016: 716) may be useful, because even though failures may remain ambiguous, failure data can help us locate problematic subsystems and joint failures that warrant special attention. Based on this analysis, we are able to identify a number of ways in which information about these ambiguous system failures can improve contracting in the supply chain in general and the warranty process in particular.

The study is structured as follows. I start by positioning my argument vis-à-vis the published literature on system failures and the managerial practice of addressing failures. There are a number of important distinctions that are required for the positioning to be clear. I then conduct an empirical analysis of system failures in the context of locomotive engines. I close by discussing the academic and practical implications.

1.2. Positioning of the Argument

It is important to position the inquiry carefully to make the intended contribution clear. Here, I position the study both academically and managerially, which I discuss in turn.

1.2.1. Theoretical Positioning

This study shares the same general research question as much of the research work on failures: How do failures occur and what can we do to improve? This said, there are a number of essential dissimilarities with published work. The main difference is that while this inquiry is about design, it is about the design of the supply chain and contractual relationships, not the product.

The product under examination—the *American Locomotive Company (ALCO) 251* diesel engine—is analogous to the Boeing 737 aircraft: a thoroughly tested, dominant design that has not changed for decades. Therefore, I assume engine design to be intact. The strongest justification for this is the fact that the ALCO 251 has been approved to be used as a standby power source in nuclear power plants (Haller, Coward, & LaMastra, 1994). It is reasonable to conclude that the design of an engine approved by the U.S. Nuclear Regulatory Commission is of sufficient quality. Further, because engine technology has been stable for decades, the major engine components across all ALCO 251 engines have an identical rated life, and both material- and design-related failures are rare. Finally, both the original equipment manufacturer (OEM) and Indian Railways have decades of experience in manufacturing and assembly. Indeed, Indian Railways assembles engines not only to the railroad industry, but also as backup sources of power at Indian nuclear power plants. Thus, this study is not about how to improve product designs through redesign, or competence in manufacturing. Instead, the focus is on implications for how to improve supply chain management in general, and buyer-supplier contracts in particular, to address these failures effectively when they do occur. In my discussion with a former CEO of an Indian Railways manufacturing facility, it became clear that engine design was hardly ever a topic when engine-related issues and problems were considered; the focus was on warranties and supply chain issues.

Second, this inquiry focuses not on how systems function but on how they dysfunction. Much of the work on complexity focuses on how the subsystems and components of a system interact in complex non-linear ways. Indeed, the classic reference in complexity theory, Simon's (1962) article *The Architecture of Complexity,* does not contain the word *failure* (or any of its variants). Of course, this is not to say that all research on complexity is about systems functioning properly (see, Agrawal, Muthulingam, & Rajapakshe, 2017; Sosa, Mihm, & Browning, 2013), but the published literature on complexity mostly examines topics such as *complex interactions* and *near-decomposability* by focusing on the architecture of a properly functioning systems. This is an important distinction, because not surprisingly, all the locomotive engine experts with whom I interacted knew the structure, even the bill of materials, of the ALCO 251 inside out. These experts also knew exactly how it was supposed to function. However, in stark contrast, these very same experts were perplexed even at the descriptive statistics of the various failure modes. For example, suppose the cylinder liner fails: will other subsystems be damaged? Which ones? How does the failure propagate? Addressing these questions must start at examining how systems fail. Here, one must distinguish between subsystem interactions in a properly functioning versus failing systems—the two are not the same thing. My focus is on the latter interactions and on the pathology of engine failures. This is analogous with the process and the purpose of autopsies: the aim is not to uncover how a person lived but how they died. Indeed, industry experts sometimes use the word "post mortem" in reference to post-failure engine teardown analyses.

I make the distinction between *architectural complexity* and *failure complexity* to make the distinction between Simon's (1962) notion of complexity and my approach clear. Architectural (Simonian) complexity is about the complexity of the product design and factors such as design specifications, bill of materials, and the number of parts. Indeed, number of parts is an often-used operationalization of complexity (Gupta & Krishnan, 1999; Novak & Eppinger, 2001; Salvador, Chandrasekaran, & Sohail, 2014). Architectural complexity can shed at least some light on how the system might fail: Where are the critical tolerances? Which parts of the system contain built-in redundancies? Which subsystems are central in the sense that their failure can induce a system failure? However, in my inquiry, the focus is on failure complexity, that is, not what we think *may happen* but what *actually happens* when the technical system fails. When a technical system fails, it functions in a way it is not designed to function. Therefore, in seeking an understanding of failure complexity, I must shift focus from the design and architecture of the product to exploring what happens when the system manifests unintended behaviors. My position is that analyzing the intended functioning of the system gives us little, if any, insight into the unintended outcomes from the system. Understanding unintended consequences is a crucial part of understanding system behavior and being able to seek improvement (e.g., Groop, Ketokivi, Gupta, & Holmström, 2017).

My emphasis is further on how *entire systems* fail. Every engine failure involves the failure of one or several subsystems, and it would be possible, perhaps even tempting, to focus on the failure of a single subsystem (cf. Ramdas & Randall, 2008). However, my research interest is at the engine level. The main motivation for this is empirical: in the 172 failures in the dataset, two thirds involve the simultaneous failure of at least two engine subsystems. This multiple subsystem failure suggests that there are potentially complex interactions associated with engine failures, and it will be both academically and managerially useful to examine what these interactions might be. When the focus is on individual component failures, multiple component failures constitute a nuisance and a threat to the validity of one’s inferences (Ramdas & Randall, 2008: 933). Multiple component failures are not a *nuisance*; they are the *essence* for gaining profound knowledge.

1.2.2. Managerial Positioning

I have worked for several years in locomotive engine maintenance, troubleshooting, and warranty fulfillment, and specifically with diesel engines similar to the ALCO 251 engine. Based on this experience, I conclude that the way failures are approached in practice is strikingly similar to the academic approach. The key principle is to focus on the failure of individual subsystems. Indeed, if a system failure involves the failure of multiple subsystems, these subsystem failures are treated as independent failures. This approach is conspicuous in the warrant claims process: if a premature engine failure involves the failure of subsystems A and B, a warranty claim is made to both suppliers A and B. Examinations of how and why the subsystem failures might be related is seldom performed. This approach may be organizationally expedient but glosses over the systematic aspects of system failures. I will show in my empirical analysis that the assumption of independent failures simply does not hold: I find systematic ways in which multiple subsystems fail. Herein resides another organizational practice with which my approach differs: in engine failure analyses, the focus is always on a single engine. In my discussions with industry experts, I did not come across attempts at analyzing failure patterns. Therefore, to my knowledge, no one in the railroad industry has conducted the kinds of analyses I have conducted. In part, the reason is that due to the contractual processes, no single vendor tends to have access to the data required to conduct such an analysis. Railroads, who do possess the complete dataset, have little incentive to engage in the analysis, their primary concern is practical: to obtain warranty replacements as quickly as possible to minimize unplanned locomotive downtime. Accordingly, many of the principles applied in the warranty process are matters of organizational expediency (cf. MacDuffie, 1997). Experienced practitioners do accumulate expertise over time, but the generalizations that individuals make based on experience are known to be highly selective and biased: in conducting a failure analysis, industry experts are more likely to recall and draw upon past failure analyses that are more recent and that were readily understandable to be recovered from memory.

Finally, it is critical to maintain the distinction between failure and defect throughout my analyses and discussions. There is indeed a tendency in the industry to equate failure with defect. An emphasis on this distinction is not to suggest that industry experts do not know the difference between the two. However, failure and defect become confounded in practice because failed components are treated in the warranty processes *as if they were* defective; again, a factual premise to be used in decision-making, not an assertion of fact (see Simon, 1997: 69, for the distinction). In this study, I challenge this factual premise and replace it with a more detailed analysis. Moving beyond this premise has implications for supply chain management as well.

In summary, my focus on the operational aspects of system failures links my research directly to operations management, and the management of operations technology in particular. In contrast with the extant literature, I do not focus on the implications for product design, but rather on supply chain design. This focus positions my study in both operations management research (Anand & Gray, 2017; Choi & Krause, 2006; Grover & Malhotra, 2003; Park & Ro, 2011; Wagner & Bode, 2014) and management research more generally (Argyres & Bigelow, 2007; Novak & Eppinger, 2001; Novak & Stern, 2008; Sanchez & Mahoney, 1996; Ulrich & Ellison, 2005; Williamson, 1971) which focuses on supply-chain management and vertical integration. In my context, because failures may have system-level characteristics, it is important to explain the ramifications of outsourcing subsystems versus making them in house.

1.2.3. The Problem with Root Causes: Architectural versus Failure Complexity

The analysis of system failures in operations management often involves the analysis of *root causes*, defined as “that most basic reason for an undesirable condition or problem which, if eliminated or corrected, would have prevented it from existing or occurring” (Wilson, Dell, &

Anderson, 1993: 3). Prescriptively, "[r]oot causes of these problems must be clearly identified and properly corrected if any real improvement is to be expected" (Wilson *et al*., 1993: 3). In examining system failures, attention typically focuses on the design of the system, and in particular, the components that failed. Post-failure analyses often uncover the failed components, but there are two complicating factors hindering identification of the root cause of failure. One factor is that it may be difficult to determine why the specific component, or components, failed. Moreover, as the quote from Mr. Belisle indicates, in the case of locomotive engine failures, a common outcome is that in many failures, no defects are found and the root cause remains elusive. What if a teardown analysis of a failed system uncovers three failed subsystems but no defects are found?

The other complicating factor is that in the case of multiple subsystem failures, it is often impossible to determine the sequence of failures (Agrawal *et al*., 2017). Did the failure of A induce the failure of B, or vice versa? Alternatively, was there a third factor that contributed to the failure of both A and B? Was this third factor endogenous (e.g., subsystem C) or exogenous (e.g., a foreign substance that entered the system and damaged both A and B)? Those conducting post-failure analyses and those seeking to improve product designs and architectures must address these questions. Often, these questions turn out to be much more complex than they think. Echoing this sentiment, Wilson *et al*. (1993: 10) noted that root-cause analysis efforts often fail, because distinguishing between apparent causes and root causes is difficult.

Elusive or not, failures link to the complexity of the system. In operations management research in particular and management and technology research more generally, scholars have focused on the design complexity of systems, or what Jacobs and Morgan (2011) called *product* *architectural complexity*. As product architectural complexity increases, the likelihood that the system exhibits emergent and unpredictable behaviors increases as well (Calvano & John, 2004; Oliver, Calvard, & Potočnik, 2017; Waldrop, 1993). However, the converse does not hold: a system need not be architecturally complex to exhibit high failure complexity. To clarify, high failure complexity does not imply that the root cause does not exist; the point is epistemological (what we can know), not ontological (how things are). Difficulties in identifying the root cause stem from the inability to observe the process of failure, only outcomes are observable. Again, autopsy—particularly in cases where the cause of death was non-natural—is a good analogy. Therefore, when I speak of ambiguous failures, I do not mean to imply that the system behaves in an ambiguous way, but rather, that understanding is ambiguous.

1.3. The Empirical Context and Analysis

Indian Railways transports daily 23 million passengers and 3 million tons of freight in its 12,500 trains across 70,000 miles of track. Between 2007 and 2015, there were 172 system failures in its fleet of 919 diesel-electric *Locomotive Engine Company* (a pseudonym) model engines. System failure in this context means that something happened in the engine that led to an unplanned shutdown and brought the locomotive to an unscheduled halt. This failure does not necessarily mean the entire system was damaged; it could be just one of the subsystems. What is relevant—and what made each of the 172 incidents enter the sample—was that the damage was sufficient to induce a system failure. Understanding how the incidents enter the sample is an important aspect of failure research (Ramdas & Randall, 2008: 931). The statistical samples of failures are of course not random, because an observation enters the sample only when a failure occurs. All statistical inferences are therefore conditional on the fact that a system failure occurred.

At the same time, because I do not focus here on explaining or predicting whether engines fail or not, but rather, how they fail. To this end, systems that do not fail provide no relevant statistical information. It is noteworthy that the system failure base-rate (172 failures in a fleet of 919 engines) is certainly high enough to warrant attention. Indeed, in my conversations with industry experts, I discerned that engine failures are surprisingly frequent and their consequences significant enough to merit attention. The annual cost of unplanned maintenance and repair and the associated engine-related warranty costs alone are in the millions of dollars. In addition, out-of-commission locomotives are a source of lost revenue for the railroad. Finally, when a locomotive engine unexpectedly fails in mid-operation, it stops all traffic for several hours. What is worse, because the majority of the railway network has only one set of tracks shared by traffic in both directions, an engine failure halts all freight and passenger traffic in all directions. This disruption also creates a potential safety hazard. Due to these ripple effects, understanding how locomotive engines fail and what can be done about these engines is a matter of urgent attention. The need for urgent attention is underscored by the annual meeting held between the railroad and all its major vendors to address persistent failures, as well as the Lac-Megantic disaster described earlier.

In sum, the railroad industry offers a potentially fruitful context for analysis of failures, because not only are the failures are frequent, public disclosure requirements give us access to the relevant data. However, examining failures in the railroad industry in particular is also societally relevant. Specifically, transportation of freight on rails is much more energy efficient and environmentally friendly than road transportation or airfreight. One rail car can carry the equivalent of two trailers, which means that a 100-car freight train corresponds to 200 semi-trailer trucks. Road transportation is not only less efficient, but also more prone to accidents, because the same road infrastructure is shared with both public and private passenger traffic. Therefore, all attempts at making rail traffic even more efficient also works toward societal ends.

1.3.1. The System under Scrutiny

The heart of the massive 132-metric-ton modern locomotive is a 15-metric-ton, 12- or 16-cylinder turbocharged diesel engine. The engine is attached to a 6-metric-ton generator, which in turn drives individual motors connected to six axles. The engine consists of power assemblies, which provide the actual power to the generator. Each power assembly consists of the cylinder head, piston, and cylinder liner, the three of which collectively connect to the camshaft and crankshaft. All these subsystems housed within the engine block that sits on the base assembly of the locomotive. Most failures are inconsequential: a faulty piston ring may cause unusual wear and tear, which does not induce a system failure, and can be addressed during scheduled maintenance operations. In contrast, a cylinder head failure is usually guaranteed to trigger a process that ultimately shuts down the engine. When the cylinder head fails in mid-operation, the engineer first notices a sharp drop in engine power. Then an alarm goes off as the Engine Protection Device activates. Depending on the nature and extent of failure, in some cases, one may notice unusually black smoke emanating from the exhaust stack on top of the locomotive (unless it is dark outside, in which case the symptom goes unnoticed). In more drastic situations, particularly those involving turbochargers, flames bellow out of the exhaust stack, and even explosions may occur. When the system fails, the engine stops, the locomotive stops, and the train stops. System failures, even the more drastic ones, are common enough to constitute not only a managerial concern but also a credible threat. This sort of failure is why, for example, “an idler car” always separates the locomotive from cars carrying hazardous cargo.

For the purposes of the analysis, I examine the engine as an aggregate of five physically distinct and architecturally decomposable subsystems: the engine block, the top deck, the cylinder head, the piston(s), the cylinder liner(s). I arrived at this partitioning based on my own engineering knowledge of the engine architecture and operation. This partitioning was further validated in consultation with locomotive experts. In this context, failure of any combination of these five subsystems can induce system failure, making the theoretical total of failure modes $2^5 - 1 = 31$. For the purposes of the inquiry, I categorize the 10 subsystem interfaces using the coupled-decoupled categorization. I define two subsystems as coupled if there is a physical interface that directly connects the two to one another; absence of a physical interface means the two are decoupled. My premise is that an understanding of how the engine can fail as a system requires and understanding of its system structure; focusing on the failure of an individual component—essentially abstracting out the system—leads to limited insight. The resulting system architecture is illustrated in Figure 1.1.

****

Insert Figure 1.1 here

****

My conceptualization of coupling is distinct from but related to definitions in the published literature. It is distinct, because in my context, the focus is on understanding failure. Because system failures may involve the propagation of the failure from one subsystem to another, I am mainly interested in which subsystems are in physical contact with one another. Accordingly, I define as coupled those subsystems that have a physical interface. In contrast, the published literature has approached coupling from the point of view of product design and redesign: two subsystems are defined as coupled if a change in the design of one has implications for the design of other subsystems (Novak & Eppinger, 2001). This means that modular product designs tend to be decoupled (Baldwin & Clark, 2000; Ulrich, 1995). However, I maintain that thinking of the engine as a system of decoupled subsystems may mislead us into thinking that system failures are subsystem specific, isolated from the broader system. However, as I demonstrate in my empirical context, the failure of one subsystem is likely to be associated with the failure of at least one other subsystem. Whether failures propagate from one subsystem to another is unrelated to whether the system is architecturally decomposable or not. Modular product architecture of course provides advantages in the maintenance and repair of the system. However, even though engines are perfectly decomposable in that they can be completely disassembled to repair and to replace components, failures often propagate from one part of the engine to another when the engine is in operation.

Sampling bias is often a concern in empirical examinations of failures (Dhanorkar, Siemsen, & Linderman, 2017; Ramdas & Randall, 2008). In my dataset, this concern is minimal, because the vast majority of system failures in the fleet of 919 engines are in the database. Importantly, the dataset contains *every* failure that occurred before the first maintenance event, which is at 365 days. The only failures missing are a handful of failures where time to failure exceeded 24 months from the date of supply of components used in assembly (the time from supply to assembly is usually a few months). The number of failures that occur after 24 months is small, but also, early failures are most relevant to the inquiry, because they are more costly and have implications for warranty. Early failures are also quite common: there are 120 failures in a fleet of

919 engines within the first year of operation before the first maintenance event, even though the rated life of the engine components under scrutiny is six years.

1.3.2. Data and Descriptives

In my data set of 172 failures, the critical data points are time to failure, and which subsystems had failed in each case. In some cases, failure of one subsystem was enough to yield a system failure; a failed piston is a good example. However, in about two thirds of the failures, it was not just one subsystem but several that had failed (Table 1.1). For each case, I also have qualitative information that gives more detailed information on the specific failure modes.

****
Insert Table 1.1 here
****

Another noteworthy descriptive statistic is that in addition to observing both single and multiple subsystem failures, I also observed three distinct failure patterns for any pair of two subsystems (Table 1.2):

1. **Concomitant failure**: the failure of one subsystem is associated with a higher failure probability in another. A potential cause of concomitant failure is the failure of one subsystem propagating to another to which it is coupled. However, it is also possible that two subsystem failures have a common cause. Two subsystem pairs exhibited concomitant failures.

2. **Preemptive failure**: the failure of one subsystem is associated with a lower probability of failure in another. Pre-emptive failure could occur as one subsystem operates as a de facto “fuse” to protect another from damaging. Preemptive failure was the most common type of relationship: six subsystem pairs exhibited preemptive failures.

3. **Independent failure**: the failure of one subsystem is not associated with the frequency of the failure of the other. Two subsystem pairs exhibited independent failures.

****
Insert Table 1.2 here
****

These failures are matters of degree. For example, no preemptive failure is such that the failure of one subsystem categorically prevents the other from failing, it merely lowers the probability. Similarly, concomitant failure does not mean that if subsystem A fails, subsystem B always fails in consequence; it merely increases the hazard. In fact, fully concomitant (B always fails when A fails) and preemptive failures (failure of A always preempts failure of B) can be thought of as the ends of a continuum with independent failures as the middle point. All actual failures in the sample fall somewhere on this continuum.

That the majority of failures involved multiple system failures merits further analysis. The likely reason for this is specifically interactive complexity (Perrow, 1984: 21-23) arising from subsystem interactions. Even though the engine is architecturally decomposable, every subsystem is linked to another subsystem: even two decoupled subsystems are ultimately coupled in the sense that there is a path of coupled relationships from any one subsystem to another. For example, the piston and the top deck are decoupled, but at the same time, the piston is coupled with the cylinder head, which in turn is coupled with the top deck. This architectural design means that concomitant failures can occur between physically decoupled subsystems as well.

1.3.3. The Hypothesis

Because failure is an unintended and unpredictable system property, it would be illogical to present specific hypotheses that predict failures. The only intellectually honest starting point is to respect our epistemological ignorance concerning highly complex systems (Dunbar & Starbuck, 2006; Hayek, 1945) and look for clues in the data. This approach echoes Hollenbeck and Wright (2017: 17) who prescribed that “[a]uthors should feel freed from the bondage of constraints imposed by a pure ratio-deductive paradigm to write more along the line of *good detective stories*.”

Indeed, trying to find out what happened and why in a system failure is analogous to a detective trying to solve a crime: the investigation is driven largely by observations; conclusions (theoretical explanations) emerge as the result of the analysis.

I shared the descriptive statistics in Tables 1.1 and 1.2 with industry experts, who found them not only non-obvious, but also puzzling in some of the observed failure modes. As an example, consider the failure mode where three physically decoupled subsystems—the top deck, a piston and the engine block—failed but the cylinder head and liner did not. How can decoupled subsystems fail if the subsystems with which they are coupled—cylinder head and liner (Figure 1.1)—do not? Where did the failure start and how did it propagate between decoupled subsystems without involving any of the coupled subsystems? This failure mode in particular echoes Mr. Belisle's sentiment that the causes of engine failures defy even systematic analysis. I also found that industry experts do not conduct systematic statistical analyses of large samples of failures. Instead, these experts apply their accumulated engineering knowledge to the evaluation of individual failures. Experts agreed that my examination of broader tendencies (such as Tables 1.1 and 1.2) is an important complement to the analysis of individual failures.

Even though I cannot present specific hypotheses, I submit that the insight into failures resides in examining subsystem interactions. I already know, from a descriptive point of view, that the majority of system failures involve the failure of more than one subsystem. However, do subsystem interactions also account for time to failure? I suggest that subsystem interactions might provide information about time to failure more than the information provided from examining single subsystem failures.

Some failures occur more rapidly. The fastest failure occurred after four days of operation, and the slowest was 1508 days, the average being 278 days. Furthermore, 71 percent of the failures occurred within one year, that is, before the first scheduled maintenance of the engines (Figure 1.2). We have essentially an Analysis of Variance (ANOVA) setup in that we have five different independent variables, one for each subsystem (Table 1.1). These variables take either a value of "0" (no failure) or "1" (failure). Of course, any combination of interactions can be modeled as well by creating interaction terms.

****

Insert Figure 1.2 here

****

There are two complicating factors, however. One factor is that because I do not manipulate the independent variables (as one would in an experimental design), the statistical effects of the independent variables on time to failure are at least partially confounded with one another. This limits the number of independent variables that can be included in the model. With five dimensions on the independent variable side, the maximum number of effects that could be modeled is as follows: main effects (five), two-way interactions (ten), three-way interactions (ten), four-way interactions (five), and five-way interactions (one). Because of the nature of the system, it is easy to see how even the five-way interaction could be significant: the idea that the way four subsystem failures interact would depend on whether the fifth subsystem failed or not is plausible. Nevertheless, we must take the confounding as given and acknowledge that we cannot isolate the interactions empirically from one another. Confounding of effects may be a cause for concern for the experimental researcher, but for us, they are simply a characteristic of the phenomenon. Figure 1.3 gives a summary of the correlations between all the main effects and the interactions, illustrating substantial confounding, which is the nature of the phenomenon. Indeed, if the effects were not confounded, they could be neatly separated from one another and estimated independently.

****
Insert Figure 1.3 here
****

As Figure 1.3 shows, modeling higher-order interactions runs into severe collinearity problems. Indeed, even trying to include several two-way interactions into the model causes estimation instability. This instability tells us that failure datasets cannot be used to conclude exactly how the subsystems interact with one another in the case of failure. For example, the two-way interaction TD*L (top deck and liner) is highly confounded with the three-way interaction TD*P*L (top deck, piston, liner); the correlation between the two is .98. There are numerous similar instances of almost full confounding. To compromise, I focus my analysis only on the *two-way* interactions, but with the caveat that the two-way interactions may well also be manifestations of higher-order interactions with which they are confounded. Trying to model higher than second-order interactions simply results in estimation difficulties. It would not make statistical sense to include three-way interactions in the model but to leave out two-way interactions (much like main effects must not be left out if there are two-way interactions).

The second complication is that the dependent variable is non-trivial in that there is a maintenance event at 365 days for each engine; a failure that occurred at 330 days cannot be compared to one that occurred at 380 days—the difference of 50 days is meaningless. In order to incorporate the maintenance event, I formulated a Cox regression model (Cox, 1972) that models time to failure using proportional hazards. In the model, I define as “failure” only those failures that occurred before the maintenance event; those that do not fail before the first maintenance event are considered "survivors," because they survive to the first maintenance event. This way the maintenance intervention is taken into consideration. The Cox model thus predicts time to failure as a function of the failure mode. The failure modes are modeled using binary variables for all five subsystems as independent variables (0=no failure, 1= failure). The Cox model is therefore essentially a Cox ANOVA with nominal dichotomous covariates, where the hazard (h) is determined as follows (e.g., Hosmer & Lemeshow, 1999: 115):

$$h(t, \mathbf{X}, \boldsymbol{\beta}) = h_0(t)*\exp(\mathbf{X}\boldsymbol{\beta})$$

where **X** is the vector of binary covariates, $\boldsymbol{\beta}$ is the vector of model parameters. Since I am interested in hazard ratios, the base hazard (h_0) is not of interest. Because the regressors are binary, the model coefficients can readily be interpreted as log-hazard ratios (Hosmer & Lemeshow, 1999: 115). To see why this is the case, consider a case with just one X that takes on either a value of 0 or 1. Now, the hazard ratio (HR) of the two is simply

$$\mathrm{HR} = \frac{h(t, 1, \beta)}{h(t, 0, \beta)} = \frac{h_0(t) * \exp(\beta)}{h_0(t) * \exp(0)} = \exp(\beta)$$

The value of the Cox model parameter β is therefore simply the logarithm of the hazard ratio. As we will later observe, this conversion makes the interpretation of the coefficients straightforward. An additional feature of the Cox model is that it does not assume any specific distribution of failure times (e.g., Melnyk, Pagell, Jorae, & Sharpe, 1995). Finally, several factors alleviate concerns of unobserved heterogeneity. For example, all engines are technologically identical and are operated in a similar operating environment by the same railroad (Indian Railways). Further, human errors are mitigated by the fact that when the engine is in operation, human discretion is highly limited because engine operations are largely computer controlled. While I cannot conclusively rule out confounding factors, it is difficult to think of ones that would have a *systematic* effect on failure times. The engines are supplied by 17 different vendors, but an analysis of failure times and failure modes reveals no link between manufacturer and failures. This lack of a linkage is plausibly explained by the fact that all 17 vendors are experienced manufacturers of the dominant-design engine.

The models were estimated using Stata 15.1.I used the original Cox specification with a non-parametric survival variable where the estimation algorithm estimates the baseline hazard as a step function that is constant between two consecutive event (failure) times. I used the `estat phtest` post-estimation procedure in Stata to test the proportional hazards assumption made in the Cox model.

1.3.4. The Statistical Results

One way to test the explanatory power of failure complexity is to compare the results of two Cox models: one that contains only main effects and another that contains (two-way) interactions. If my general conjecture holds, the model with the interactions should have more statistical explanatory power, which, of course, also implies establishing root cause is arduous. For example, suppose we take a conventional, component-centric approach, and focus only on cylinder head failures and compare failures where the cylinder head failed to failures where it did not fail. Does this piece of information on one subsystem failure give us any additional information on whether the failure was more likely to be above or below the sample average of 278 days? My conjecture is that a single subsystem failure gives less information, because failure time likely depends on whether other subsystems failed in addition. Simply put, isolating a single subsystem for analysis may be analytically convenient but ultimately does not do justice to how these systems fail (cf. Ramdas & Randall, 2008).

In Model 1 (Table 1.3), I have modeled all five main effects, and the model for the hazard is therefore

$$h(t, \mathbf{X}, \boldsymbol{\beta}) = h_0(t)*\exp(\beta_{TD}X_{TD} + \beta_{CH}X_{CH} + \beta_P X_P + \beta_L X_L + \beta_{EB}X_{EB})$$

The variables X_{TD}, X_{CH}, X_P, X_L, and X_{EB} take on values 0 or 1, depending on which subsystems fail. The assumption in this model is that individual components have additive predictive effects on the system failure hazard. The model parameters are reported in Table 3.

****

Insert Table 1.3 here

****

As we can see from Table 1.3, none of the betas are significant, which means that there is no evidence of independent, additive effects predicting failure. Substantively, this means that a component-centric analysis does not provide *any* insight to time to failure. In the other models, we proceeded to examine the predictive power of the two-by-two interactions. To this end, I ran models with all 10 two-way interactions and found five interactions to be significant (Models 2 through 6). The proportional hazards assumption seems to hold in all models reported in Table 1.3 (the chi-squares are insignificant for all models). Just to be sure, in addition to the global test for each model (with df equal to the number of covariates), I also ran the test at the level of individual covariates in each model. None of the covariate-specific chi-squares (with 1 df) were significant in any of the models, therefore, I am confident that the proportional hazards assumption is not violated. In the following, I discuss the significant interactions.

1.3.5. Meaningful Interpretations through Explanatory Induction

The Cox regression models give us a reason to believe that time to failure links not to whether and how individual subsystems fail, but rather, how subsystems interact in the case of failure. I had no *a priori* expectation as to how exactly these interactions would occur, but the Cox models provide us with important clues. To be clear, I do not suggest that the significant interactions test specific hypotheses about subsystem interactions or the causal processes associated with the failures; I have only observed the outcome of the failure. Subsequently, I have examined these outcomes through statistical analysis to uncover tendencies that merit further analysis. In methodology literature, this approach is sometimes referred to as *explanatory induction* (Rozeboom, 1997: 381) as opposed to the more familiar *hypothetico-deductivism* (Hempel, 1965). The hypothetico-deductive research design is generally useful when theory is strong enough to yield predictions about specific model parameters. However, in the context of complex failures, such specific predictions are impossible to make; indeed, the central managerial challenge is precisely *not* knowing how the engines fail and what the dominant interactions might be.

That the main-effects-only model had no predictive power but the interaction models did is my primary clue, which requires further examination and interpretation. Without plausible interpretations, the mere statistical significance of these interactions does not supply meaning. In the words of Rozeboom, we must turn “observed pattern parameters into stories about what seems to be responsible for what” (1997: 381). I addressed this concern by presenting the findings to industry experts and asking them to think of plausible interpretations. The following interpretations are based on these conversations. I discuss two of the five observed interactions as examples.

1.3.5.1. Why Does Joint Failure Predict Faster Failure (Model 5)?

Consider first the interaction of the piston and the engine block (Model 5). The hazard ratio is 4.51, which means the joint failure of the two is associated with over a four-fold failure hazard, and therefore, predicts a significantly shorter time to failure. To arrive at a meaningful interpretation of this empirical finding, we must go back to the system structure (Figure 1.1). The engine block and the piston are decoupled. Jointly, they exhibit a preemptive failure tendency (Table 1.2), which is to be expected for decoupled subsystems: if one subsystem fails, the entire system fails before any decoupled subsystems are damaged. However, here is the puzzle: Why does the joint failure of the two link to a *shorter* time to failure? This is where the descriptive statistics (Table 1.1) can be useful. Specifically, in the failures where both the engine block and the piston failed, other subsystems tended to fail as well; indeed, in the majority of cases it was at least *four* subsystems that failed. The joint failures of these two decoupled subsystems seem to be associated with massive failures that damage many other subsystems as well. The next question to be addressed is why these massive failures would occur, on average, earlier than other failures. The plausible explanation is that there was something wrong with the system at the inception. Potential causes for this failure are faulty components or faulty assembly.

The first-order coefficient of the engine block in Model 5 is below one, indicating a lower failure hazard. This empirical finding suggests that if the engine block fails but a piston does not (making the interaction term zero), it takes longer for the engine to fail. Going back to the descriptives, Table 1.1 indicates that in the vast majority of cases where the engine block failed but a piston did not, no other subsystem failed. Therefore, the event is a single-subsystem failure.

Further, the single subsystem that fails is the four-ton engine block. It understandably takes more time for this massive part of the engine to fail.

1.3.5.2. Why Does Joint Failure Predict Slower Failure (Model 4)?

Positive interactions between two subsystems (i.e., positive log-hazard ratios or hazard ratios greater than one) means that a joint failure of two subsystems is associated with more rapid system failures. However, we also observed two negative interactions (i.e., negative log-hazard ratios or hazard ratios less than one), one of them being the interaction of the piston and the cylinder liner (Model 4); the joint failure of these two subsystems is associated with a longer time to failure. The liner is a hollow cylindrical sleeve that houses the piston. The two subsystems are not only coupled (Figure 1.1), but also when the engine is in operation, the piston moves cyclically up and down inside the liner sleeve at a rate of 400-1000 times per minute, making the two also *dynamically* coupled. This subsystem contrasts with *statically* coupled subsystems, such as the cylinder head and the engine block. The predictable failure for the subsystem of the piston and liner mode is wear and tear that damages them both. Indeed, the two tend to exhibit concomitant failure (Table 1.2). Because wear and tear takes time, time to failure can be expected to be longer. Given this time lag, when a tightly coupled subsystem fails, for practical purposes, it still can function like a loosely-coupled subsystem, at least for a time, which is a real-world characteristic of subsystems that the language in the extant literature typically does not capture.

The first-order coefficients for piston and liner in Model 4 are both greater than one. Again, just as in Model 5, these are the conditional main effects of one subsystem failure given the other subsystem did not fail. These conditional main effects that exceed one mean that in the event that one of the subsystems fails but the other does not, the engine fails faster. However, how is it possible for the piston to be damaged but not the dynamically coupled sleeve that houses it, or vice versa? The plausible explanation is another subsystem failure that damaged either the piston or the liner, but the system failed before the failure in one propagated to the other. Table 1.1 offers clues to examine what the other source could be: in all the cases where the piston failed but the cylinder liner did not, either the top deck or the cylinder head, or both, had failed. This finding is important, because even though the piston and the liner exhibit a concomitant failure tendency, the number of system failures where only one of these subsystems failed (47) exceeds the number of system failures where both subsystems failed (34). This finding means that even in the event of a system failure, the failure of one subsystem may damage the subsystems to which it is coupled, but that it is more likely that this outcome does *not* occur. A contribution of this study to the extant literature is the recognition that an architecturally tightly coupled subsystem can vary significantly *in degree* in terms of the time it takes a failure to propagate through the subsystem.

Regardless of the degree to which failures propagate within a subsystem, I have established empirically that failures propagate across subsystems, which implies that locomotive engines, in practice, are not as “nearly decomposable” as some might imagine. This lack of decomposability poses especially challenging contractual problems for fragmented supply chains because no individual decision maker has the incentive to identify the root causes of system failure, which thus impedes learning. The next section considers how the intrinsic challenges inherent in complex technological systems and a variety of contractual arrangements co-evolve.

1.4. Discussion

Locomotive engine failures are not exceptional events. Indeed, they are so common that Indian Railways must address them as a matter of their ongoing organizational concerns. When I queried top management and engineers at the manufacturing facility regarding what they considered their most vexing challenge, their response was precisely the management of the supply chain and quality. For Indian Railways, the fragmented value chain further complicates matters: not only does Indian Railways outsource the engine design and production to engine manufacturers, but also these manufacturers further outsource the production of engine subsystems to specialized manufacturers. How are engine failures handled in this fragmented supply chain? What are the principles and policies by which costs of failure are assigned in failures that involve multiple engine subsystems, particularly in situations where the root cause of the problem cannot be identified? In this discussion section, I focus on the organizational implications of the empirical findings.

To set the stage for examining the organizational challenge, I first describe the inter-organizational value chain of the locomotive engine and the warranty process, and in particular, Indian Railways' role in it. The key actors of the value chain are depicted in Figure 1.4. Indian Railways performs the final assembly of the engine using subsystems supplied by Locomotive Technology Company (LTC, a pseudonym for the OEM supplier). LTC is an engine technology specialist with competencies in engine and subsystem design. Indian Railways participates in the engine design process through its internal research division, which provides the general performance specifications—"the performance envelope"—for the engines; the actual engine design is, however, LTC's responsibility. Assigning the responsibility of overall system and subsystem design to just one organization is crucial, because changes in environmental standards, for example, may require both system- and subsystem- level modifications to the engine. Having individual suppliers design their own subsystems would lead to high coordination costs. Again, in the case of the ALCO 251 engine, design changes have been minimal and further redesign is not likely. The central inter-organizational challenges are not technical, but rather managerial and contractual.

****

Insert Figure 1.4 here

****

In addition to using engine suppliers such as LTC, Indian Railways also considers both assembly and disassembly of engines a central internal competence. Being able to assemble and disassemble engines in house expedites engine maintenance and repair and ensures that Indian Railways is not completely dependent on engine suppliers. When Indian Railways performs manufacturing activities in house, its Engine Assembly Unit orders both new and replacement subsystems from LTC based on its own assembly and operational requirements. More generally, external suppliers provided all replacement parts and subsystems. Understandably, it would make little economic sense for Indian Railways to manufacture, say, its own liners. Instead, subsystem manufacturing is best left for specialized manufacturing firms.

When an engine failure occurs, the locomotive is first hauled to one of Indian Railways' roughly fifty engine maintenance depots. Technical experts at the depot perform a routine tear-down analysis to determine which subsystems have failed and been damaged. If deemed appropriate, the depot maintenance facility then submits a warranty claim and a request for replacement subsystems to a vendor (e.g., LTC). The Engine Assembly Unit relays the depot's request to the vendors. If the vendors grant the warranty, the vendor supplies a replacement subsystem directly to the engine maintenance depot; if the warranty is denied, Indian Railways purchases the replacement subsystem. Importantly, warranty claims are made only for individual subsystems: when the depot submits a warranty claim for, say, a cylinder liner, all information on whether the failure had involved other subsystem failures is ignored. If three different subsystems have failed, three separate warranty claims are made independently of one another, one for each subsystem. Therefore, for all practical purposes, and as a matter of corporate policy, every failed subsystem is in some sense treated as *the* root cause. In my dataset, the 172 failures were associated with 334 separate warranty actions. According to my conversations with industry experts, this component-centric approach to warranties is standard industry practice throughout the world, and is not limited to the railroad industry (e.g., Agrawal *et al.* 2017).

The engine maintenance depots are responsible for minimizing unplanned locomotive downtime. Their incentive is therefore to submit subsystem warranty claims as quickly as possible in order to receive replacement parts swiftly. This process is often expedited by the engine assembly unit, which supplies the depot with replacement subsystems when making the warranty claim. Since the engine will eventually need the replacement part anyway; why not supply it immediately from the internal spare-parts inventory? The objective is to return the failed locomotive back to service as soon as possible. For the most part, failed subsystems are replaced at the depots; only the most challenging cases are transferred to the Engine Assembly Unit. Prima facie, the component-centric approach makes perfect sense when the overriding objective is to reduce unplanned downtime.

Vendors who respond swiftly to warranty claims are preferred, even to the point that high responsiveness may offset the negative effect the frequency of failures has on their performance evaluation. However, the problem is that while the swift, component-centric warranty process may work towards minimizing locomotive downtime, it may also effectively disguise the underpinning problems precisely because it disregards all systematic drivers of failure. Again, two thirds of the failures in the dataset involved two or more subsystems failing simultaneously. In the current warranty process, the primary objective is to replace failed parts as quickly as possible, there is little incentive to examine and understand system-level aspects or complex interactions associated with the failures.

Quality control practices at Indian Railways are also component centric in that failure statistics are tracked and analyzed *at the component level*. This practice is precisely what led to the downgrading of the subsystem supplier in the example described in the Introduction. When a maintenance depot observed an unusual frequency of failures in the subsystem, it reported the finding internally to the Engine Assembly Unit and the research division and externally to the subsystem vendor. The internal team further collected failure data on the subsystem from all depots across India. All attention was focused on the specific subsystem and its failure rate, ultimately leading to the mistaken conclusion that the vendor was at fault. The assumption is that unless there is strong evidence to the contrary, the failed subsystem is considered to be the root cause: it failed because there was something wrong with it.

Let us examine in more detail how the component-centric approach may create distortions; cylinder liner failures are a good example. The cylinder liner failed in 46 system failures, but almost 80 percent of these failures involved the failure of at least one other subsystem (Table 1.1). Further, the piston and the liner also exhibit a strong concomitant failure pattern (Table 1.2): in the absence of a piston failure, the probability of a liner failure before the first maintenance event is about 12 percent, but with a piston failure, the probability of liner failure is *four-fold*, 49 percent. Finally, the interaction term of the liner and the engine block is in the Cox model indicates that when the two fail simultaneously, the failure is more likely to occur earlier, before the first maintenance event. Ignoring these conspicuous system-level phenomena may lead to an unfair allocation of responsibility for system failures. At the same time, these distortions are remediable by giving up the exclusively component-centric approach in favor of at least a partial system-centric approach. In the following sections, I examine this approach in more detail in light of our findings on joint failures of multiple subsystems in particular.

1.4.1. Current Managerial Responses

Figure 1.5 gives a summary of the interactive complexity in engine failures. The horizontal axis places the joint failure of the two subsystems on the preemptive-concomitant continuum. The x-coordinates are calculated from Table 1.2 as conventional log odds ratios:

$$\mathrm{x} = \log \frac{\mathrm{P(B\ failed\ |\ A\ failed)}}{\mathrm{P(B\ failed\ |\ A\ did\ not\ fail)}}$$

Concomitant failures have positive values, preemptive failures negative values, and independent failures zero values on this continuum; log odds are used precisely because it gives independent failures a value of zero. The vertical axis indicates whether the joint failure of the two is associated with faster or slower time to failure. The y-coordinate is the logarithm of the estimate of the interaction from the Cox model. Faster failures have positive and slower failures negative values on the vertical axis— a value of zero indicates no relationship. Finally, frequency of failures merits attention: there were 58 joint failures involving the top deck and the cylinder head, but only 4 that involved the cylinder head and the engine block. The area of each circle is proportional to the number of joint failures observed in the data. Figure 1.5 gives a succinct summary of the nature and the magnitude of the operational (interactive) complexity challenge for each subsystem pair.

****

Insert Figure 1.5 here

****

The component-centric warranty process would work well only if the observations in Figure 1.5 were all close to the origin. For instance, the top deck and the cylinder liner fail more or less independently of one another and the joint failure of the two neither accelerates nor decelerates time to failure. Their joint failure presents no challenges to the warranty process or management of the vendors and supply chain. If the two fail jointly, the warranty claims for the two subsystems can be directed at the manufacturers of the two subsystems, whether they are the same firm or two different firms.

However, the majority of subsystem pairs in Figure 1.5 are far from the origin. The joint failure of the piston and the cylinder liner is the best example. Not only are joint failures of the two common (34 of all failures involve joint failure of the two), the two also exhibit the strongest concomitant failure pattern we observed in the data: the failure of one of the two increases the odds of the other failing by a factor of *seven*. This is a clear indication that the failures of the two are connected. However, establishing the root cause is very difficult: it may be simply impossible to establish whether piston failure was caused by liner failure, or vice versa. What is more, in cases where a pair of subsystems failed, it is possible—even likely—that at least a third subsystem failed as well. This failure pattern means that concomitant failure of two subsystems may be spurious as well.

Even a component-centric approach must incorporate the joint failures in one way or another. Currently, Indian Railways' policy is not to hold a subsystem vendor liable if the damage to its subsystem was unambiguously caused by the failure of another vendor's subsystem.

Similarly, a subsystem vendor will not be held liable if its subsystem caused other subsystems to fail. This is standard industry practice, because no subsystem vendor would accept the risk of being held responsible for such consequential damage. If they had to bear such risk, they must carry expensive liability insurance policies, which would predictably drive up the cost (and the price) of the subsystem. However, the main problem with the component-centric approach stems from the fact that the majority of the failures involve several subsystem failures that may be connected to one another but where unambiguously establishing the root cause may be impossible. The problem is further exacerbated by the fact that subsystems are often provided by different vendors. If the manufacturers or subsystems A and B were the same firm, it would be just an organizational problem for the joint manufacturer of the two subsystems. However, because the manufacturers are often not only separate firms but also entities separate from Indian Railways, assigning responsibility is complicated. Understandably, if consequential failure of subsystem B due to the failure of subsystem A is suspected but not firmly established, the supplier of subsystem B may challenge the decision of being held responsible, particularly if it has provided one of the more expensive subsystems. Industry experts confirm that this is an authentic problem: vendors often disagree on how liability should be assigned, and sometimes disagreements result in arbitration or even litigation. In the case of concomitant failure, root cause ambiguity arising from inter-dependencies is usually a source of delays and friction and enforcing the vendor contracts is far from trivial, nor can it be addressed by modification of the component design (cf. Agrawal *et al*. 2017; Sosa *et al*. 2013). Finally, all conflicts are a cause for concern in a context where the different actors in the value chain are fundamentally dependent on one another and where supplier switching-costs are considerable (cf. Monteverde & Teece, 1982). Currently, almost one in seven suppliers to Indian Railways are blacklisted as a result of disputes, often stemming from quality-related issues.

1.4.2 Rethinking Contract Design and the Warranty Process

A single vendor supplying all subsystems would eliminate the problem with assigning warranty liability. Indeed, some railroads adopted this approach in the 1990s through various Engine Reliability Programs, which contractually transferred responsibility for in-service availability from the railroads to the OEMs. This practice, generally known as *performance-based contracting*, is common in the transportation industry in particular (Guajardo, Cohen, & Netessine, 2016). Railroad industry experts with whom we interacted confirmed that adopting this approach not only eliminated the vast majority of disagreements, but it has also expedited the warranty process. The downside is railroads generally consider these contracts prohibitively expensive. Another downside is the loss of internal competence for the railroad: the more one outsources the more one loses in competence. Eventually, industry practice gradually shifted toward fragmentation, specialization, and the component-centric approach. This corresponds to conventional after-sales service, sometimes dubbed *time-and-material contracts* (Guajardo *et al*. 2016: 960).

The industry also experimented with a Unit Exchange Program, where subsystem vendors replace (“exchange”) individual components (“units”) as a matter of preventive maintenance; this is essentially applying the time-and-material contract. For example, a vendor replaces cylinder liners that have a rated life of six years every three years. The vendor is also responsible for their failures. The program is employed by several railroad operators and is generally recognized as efficient in terms of minimizing unplanned downtime. However, the results of this study indicate that this is problematic if only some of the subsystem suppliers participate. For example, how much sense would it make to replace, as matter of preventive maintenance, one subsystem (e.g., piston) if the same was not done to another subsystem that tends to fail concomitantly with it (e.g., liner)? Viewing the question from the component-centric perspective would simply be glossing over this problem. However, despite the problems associated with the Unit Exchange Program, some variant has been widely adopted in North America in particular. Railroads expect that the revenues lost due to downtime exceed the costs of adopting these programs by several orders of magnitude.

My analyses may offer specific guidance on how to rethink of the warranty process. The central idea is to shift towards a system-centric perspective, but in an analysis-driven, selective manner: everything hinges on understanding Figure 1.5. Specifically, some subsystems are more problematic than others are, not because they fail more often but because in case of failures they interact with other subsystems (see, Novak & Wernerfelt, 2012). Concomitant failures are particularly problematic: it may be impossible to determine whether the failure of a subsystem was the consequence of another subsystem failure or whether the subsystem was defective (Agrawal *et al*. 2017).

A closer reflection of the results reveals that most near-zero correlations involve the top deck, and once I include the engine block, we have accounted for all but one case. Further, the three subsystems that comprise the power assembly—cylinder head, piston, and liner—are associated with the most problematic interactive failure patterns. This is not surprising; after all, given the combustion cycle occurs within the power assembly. Given these technical realities, combining the three subsystems that comprise the power assembly into a single subsystem, to be provided by the same vendor, will alleviate contractual problems. The interactions of the three subsystems would still present a challenge to the power assembly supplier, but the problem would transform from a contractual problem to an internal engineering and management problem. Problems are always tackled more effectively within organizations than between organizations.

My primary managerial prescription is the following: based on my analysis, we suggest that Indian Railways buy the power assembly from vendors who supply the entire assembly and are responsible for all its components. In addition to turning a comparatively more complex inter-organizational problem into a more manageable intra-organizational problem, it would also directly reduce contractual costs: in the case that two or three subsystems within the power assembly failed, only one warranty action would be needed. Based on my data, this would reduce the total number of warranty actions by 17 percent; the number of warranty claims associated specifically with the subsystems within the power assembly would be reduced by 30 percent. These reductions in warranty claims are associated with significant reductions in administrative costs that should always be considered in conjunction with production costs. Further, these costs can unambiguously be assigned—in the spirit of activity-based costing—to specific buyer-supplier contracts, they do not (and should not) be treated as overhead.

My discussions with a former CEO of Indian Railways' manufacturing unit reveal that the recommendation is not only is it feasible, it can be implemented within the existing organizational and policy framework. Whether the new approach would make the supply chain more efficient is an open empirical question that requires a thorough evaluation. The expectation that administrative costs would decline is highly warranted, but these cost savings must be compared to potential increases in the subsystem price. What is the cost of the power assembly compared to the cost of buying it components and assembling them in house? I do not have the precise data to conduct these calculations, but the requisite data can be collected and the question can be addressed. However, I speculate that since even a few days of locomotive downtime (e.g. due to warranty disputes and non-availability of replacement parts) likely exceeds the cost differential, the modification in contract design may well be justified.

Beyond cost savings, it would be useful to analyze whether the new arrangement would have also implications for product quality. Would power assembly failures eventually decline? Ultimately, the problem should be addressed *in its entirety*, which involves both contractual and quality issues. The reason I might hypothesize a reduction in failures is that if one specialized vendor was in charge of not only power assembly but also the sole originator of warranties for all its components, this approach could create an incentive for the vendor to engage in more rigorous quality management. This approach would be further supported and enabled by the fact that the power assembly provider would get access to data that are more detailed concerning how it failed, and specifically which subsystems failed in conjunction with one another. In the current system, these crucial insights get lost in the component-centric warranty process.

1.5. Conclusion

The prescription that organizational structures and supply chain structures should mirror the technical realities of the task at hand is well established. However, this prescription hinges on the critical assumption that the relevant technical realities are known. In this dissertation chapter, I have examined a context where this is not the case. On the one hand, the product itself—the locomotive engine—is modular and, consistent with the mirroring hypothesis, the subsystems are supplied by specialized vendors. On the other hand, these engines exhibit high failure complexity—failures involve multiple subsystems and root causes are often impossible to identify.

It is further difficult to determine whether a failed subsystem was genuinely defective. This presents a challenge for the fragmented supply chain: In the case of ambiguous joint failures, which vendors will be held responsible for what?

I have sought to shed more light on unexpected failures by a systematic investigation of 172 locomotive engine failures. Even though the ambiguity of the failures cannot be resolved, the analysis does reveal that subsystems exhibit variable degrees of *failure complexity*. The general prescription is to combine the particularly problematic subsystems—however modular—into a single subsystem, supplied by the same vendor. In practice, this approach has not been implemented, instead, the supply chain mirrors the logic of modularity where architecturally decomposable subsystems are outsourced to different vendors and each vendor is responsible for the failures of the subsystems they supply. This practice is organizationally understandable but may not align with the technical realities of the failures. Decisions regarding supply chain structure in uncertain environments should not be made simply based on product architectures and designs but also by incorporating the technical realities of failures. I have shown that even though failures will probably always elude us, a systematic analysis can lead to an improved understanding. In particular, when warranty policies assign liability in a transparent and reasonable manner, everybody benefits in the long term.

FIGURE 1.1

The Architecture of the Locomotive Engine

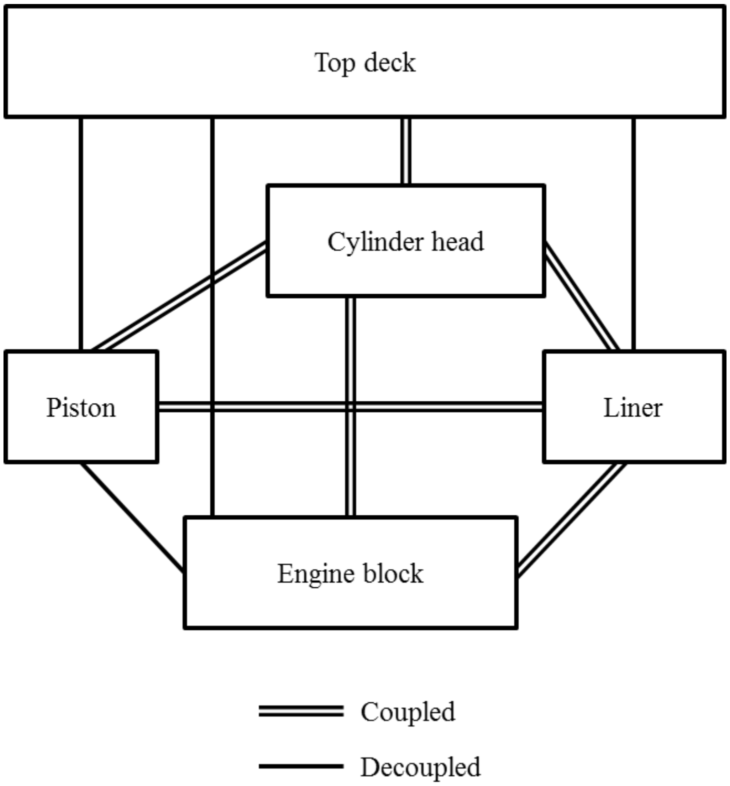

TABLE 1.1

Failure Modes and Frequencies

Top Deck	Cylinder Head	Piston	Cylinder Liner	Engine Block	Number of failures	Frequency	Percent	Cum. percent
1	1	0	0	0	multiple (2)	41	23.8	23.8
0	0	0	0	1	single	25	14.5	38.3
1	0	1	0	0	multiple (2)	14	8.1	46.5
1	0	0	0	1	multiple (2)	13	7.6	54.0
0	0	1	0	0	single	12	7.0	61.0
0	0	1	1	0	multiple (2)	11	6.4	67.4
0	1	0	0	0	single	10	5.8	73.2
0	0	0	1	0	single	9	5.2	78.5
1	0	1	1	0	multiple (3)	9	5.2	83.7
1	1	1	1	0	multiple (4)	8	4.7	88.3
1	1	1	0	0	multiple (3)	5	2.9	91.2
1	1	1	1	1	multiple (5)	3	1.7	93.0
0	1	1	0	0	multiple (2)	2	1.2	94.1
1	0	0	0	0	single	2	1.2	95.3
1	0	1	1	1	multiple (4)	2	1.2	96.5
0	0	0	1	1	multiple (2)	1	.6	97.1
0	1	0	1	0	multiple (2)	1	.6	97.6
0	1	1	1	0	multiple (3)	1	.6	98.2
1	0	0	1	1	multiple (3)	1	.6	98.8
1	0	1	0	1	multiple (3)	1	.6	99.4
1	1	1	0	1	multiple (4)	1	.6	100.0
				Total	single	58		33.7
					multiple (#)	114		66.3
				Total		172		100.0

TABLE 1.2

Two-by-Two Relationships of Subsystem Failures

		Cylinder Head		**Piston**		**Cylinder Liner**		**Engine Block**	
		No Fail	Fail	No Fail	Fail	No Fail	Fail	No Fail	Fail
Top Deck	No Fail	**58**	**14**	46	26	49	23	**46**	**26**
	Fail	**42**	**58**	57	43	77	23	**79**	**21**
	Type	Concomitant		Independent		Independent		Preemptive	
Cylinder Head	No Fail			**51**	**49**	**67**	**33**	**57**	**43**
	Fail			**52**	**20**	**59**	**13**	**68**	**4**
	Type			Preemptive		Preemptive		Preemptive	
Piston	No Fail					**91**	**12**	**63**	**40**
	Fail					**35**	**34**	**62**	**7**
	Type					Concomitant		Preemptive	
Cylinder Liner	No Fail							**86**	**40**
	Fail							**39**	**7**
	Type							Preemptive	

The two-by-two tables that appear in boldface indicate a statistically significant association at p=.05; gray shading indicates subsystems that are coupled, white cells indicate decoupling.

FIGURE 1.2

Distribution of Times to Failure

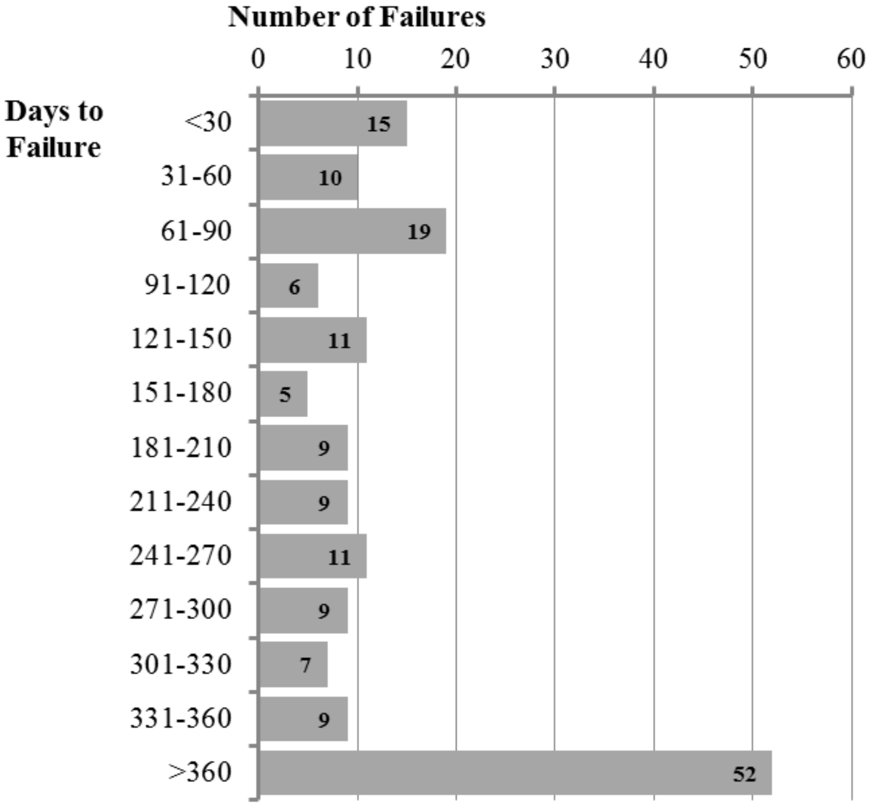

FIGURE 1.3

Confounding of the Effects

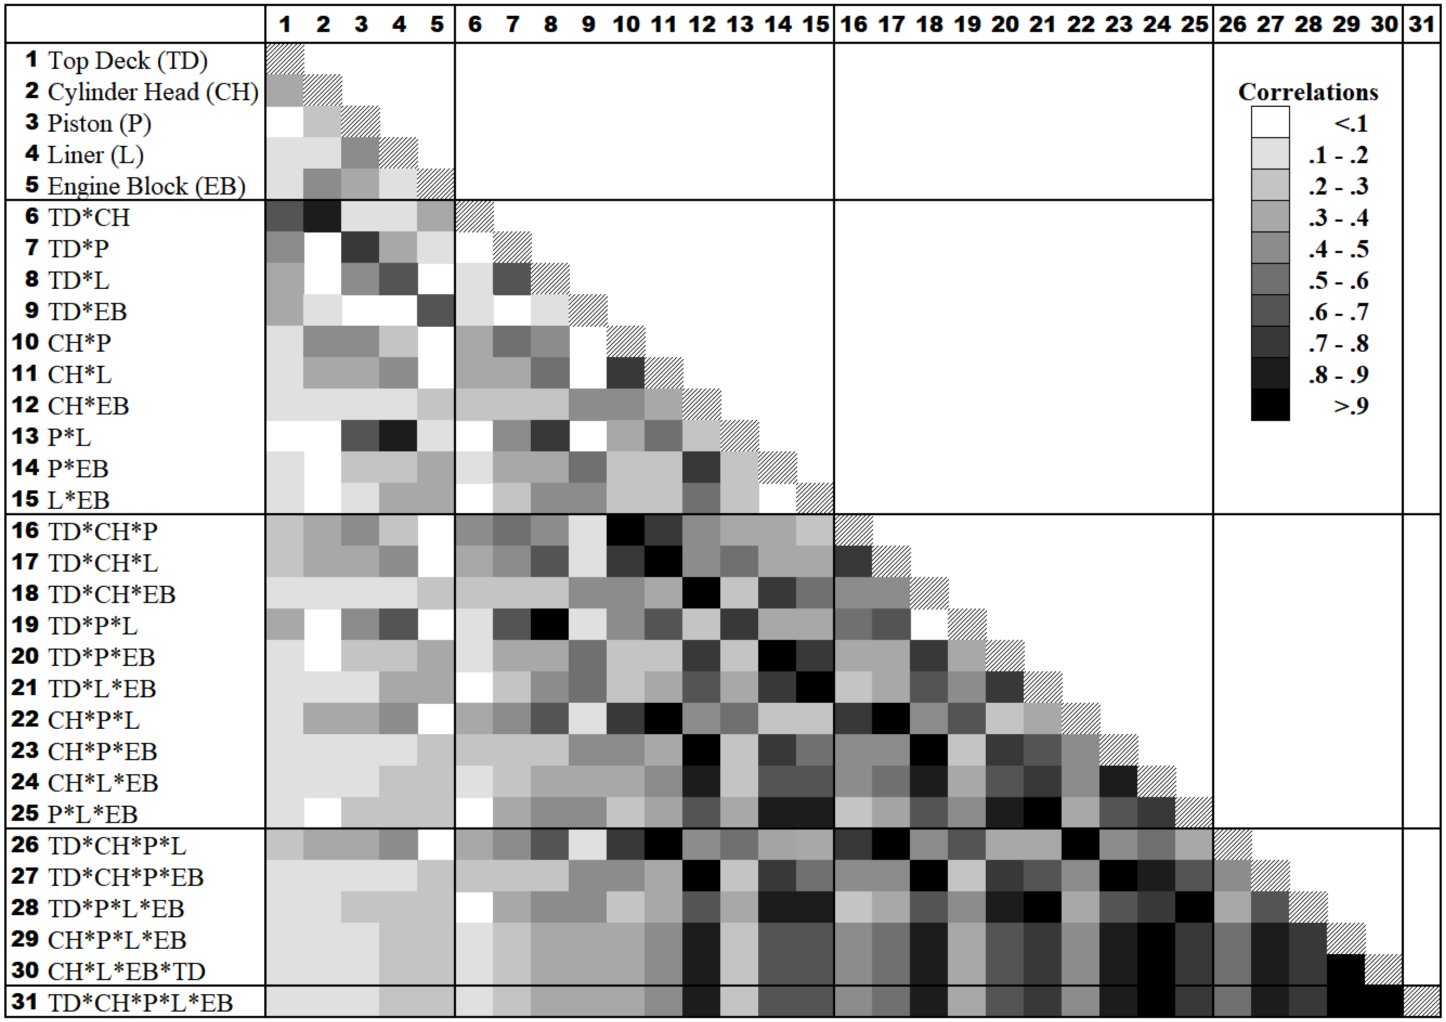

TABLE 1.3

Modeling Time to Failure: Cox Regression

	Model					
Variable	[1]	[2]	[3]	[4]	[5]	[6]
Top Deck (TD)	.99	.57 *	1.24	1.10	.94	.97
Cylinder Head (CH)	.73	.70	.72	1.00	.54 *	.59 *
Piston (P)	1.05	.53 †	1.10	2.00 *	.78	.91
Cylinder Liner (CL)	1.39	1.23	1.46	4.73 ***	1.25	1.18
Engine Block (EB)	.66	.53 *	1.00	.92	.41 **	.46 *
TD * P		3.29 **				
TD * EB			.422 *			
P * CL				.17 ***		
P * EB					4.51 *	
CL * EB						3.24 *
Model fit						
Likelihood ratio chi-square	10.3	19.7 **	14.1 *	22.1 **	15.4 *	14.0 *
df	5	6	6	6	6	6
Proportional hazards chi-square	4.67	5.23	5.43	6.15	6.02	5.71
df	5	6	6	6	6	6

† p<.10 * p<.05 ** p<.01 *** p<.001
Estimates are the hazard ratios

Note: Only significant interaction effects presented: the remaining two-way interactions were tested but were not significant.

FIGURE 1.4

The Locomotive Engine Value Chain

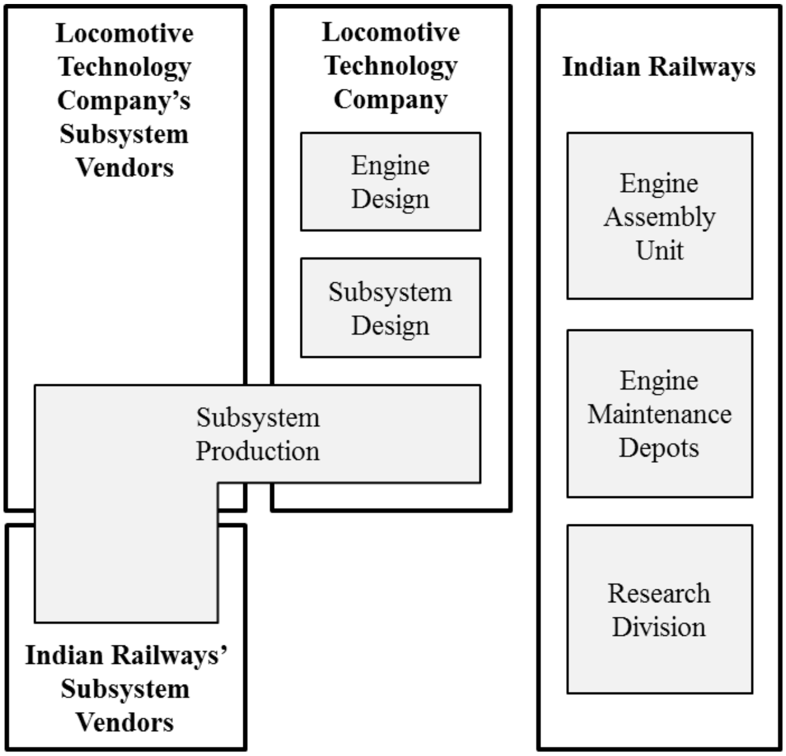

FIGURE 1.5

Failure Complexity Summarized

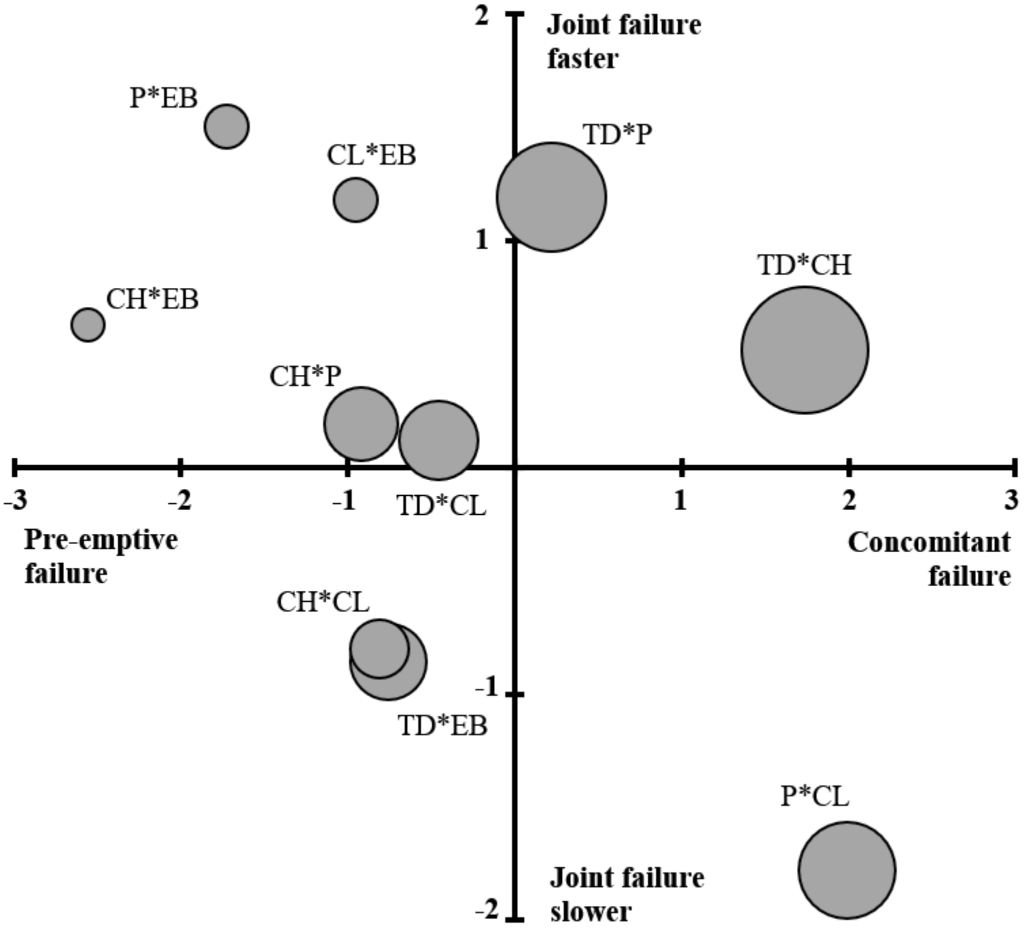

X-coordinate:	log odds of the joint failure
Y-coordinate:	log of the coefficient of the interaction term in the Cox model

CH	cylinder head
CL	cylinder liner
EB	engine block
P	piston
T	top deck

Chapter 2

Opportunism or Honest Disagreements? Evidence from Buyer-Supplier Contracts in the Railroad Industry

Abstract

This study considers contractual disagreements between Indian Railways and its vendors between 2002 and 2016. I provide four *opportunism-free* scenarios in which there are contractual disagreements or contract termination over the quality of failed components. Such disagreements occur because managers are dealing with complex technical systems, which are intrinsically fraught with ambiguities. Thus, observing (unexpected) component failures typically confront managers with ill-structured problems that lead to contractual disagreements by the two exchange parties (i.e., Indian Railways and its focal vendor of the component in question) about the causes of these technical failures. These disagreements do not yield technical resolution, but are resolved by managers acting in the 'spirit of the contract.' When managers who had not fully evaluated technical reasons for component failure, assume, without evidence, that their transacting partner is behaving opportunistically, and enforce the 'letter of the contract,' contract termination occurs. Better management practice obtain when managers adhere to the 'spirit of the contract,' and inferences of opportunism are evidence-based, rather than a default premise invoked to explain (unexpected) failures of complex technical systems.

Keywords: honest disagreements, opportunism, complexity, bounded rationality, transaction cost economics, contracting.

2.1. Introduction

In 2012, due to a sudden spate of early power assembly failures, Indian Railways downgraded a Fortune 500 original equipment manufacturer (OEM) vendor from its preferred (Tier 1) status. To do so, the railroad enforced strictly, a contractual clause that grants it discretion to 'punish' any vendors that fail to meet the contractually stipulated quality requirements. Indian Railways awarded Tier 1 status and the supply contract worth millions of dollars to an alternative Fortune 500 OEM. In part, the reason for the adopted course of action was Indian Railway's underpinning premise that the original vendor's supply was of poor quality and that this supplier in denying warranty, was acting opportunistically (Grover & Malhotra, 2003; Williamson, 1996).

Supply chain management (SCM) research which focuses on potential opportunism is likely to find the railroad's response reasonable. For example, to safeguard against the hazards of opportunism, Gray and Handley (2015) caution against using a single contractor; Terpend and Krause (2015) promote the use of competitive pressures between suppliers to improve performance; Liu, Luo, and Liu (2009) advocate using formal contractual mechanisms; Lonsdale (2001) highlights the hazards of post contractual dependency; and Zu and Kaynak (2012) suggest focusing on incentives and rewards. Surprisingly, however, the intervention proved unsuccessful as the components supplied by the other OEM vendor also exhibited a similar rate of failure.

Indian Railways is one of the largest railroads in the world with considerable expertise. It owns and operates two locomotive assembly facilities producing hundreds of locomotives annually, and all vendors undergo a stringent qualification process. Similarly, the OEMs have decades of expertise and experience in manufacturing. The technology associated with locomotive production is also relatively stable. Yet, the buyer-supplier relationship is fraught with similar problems – overall, approximately 14 percent of vendors' contracts supplying components to Indian Railway are terminated indicating severe disagreements over the quality of supplied components and technical issues. What's going on here?

It would be unwise to suggest that opportunism has no role to play in any contractual disagreements and contract terminations. In my study, I submit that earnest transacting partners can have honest disagreements, and not *all* contractual disagreements that occur are a consequence of opportunism. I posit that these contractual disagreements and contract terminations arise as result of managers' incomplete understanding of the complex technical systems they seek to manage (Ramasesh & Browning, 2014). I focus on identifying *when* and *why* honest disagreements are more likely to occur, and how these contractual disagreements and contract terminations may be attenuated.

In the interest of advancing the extant state of theorizing about economic exchange, especially in the buyer-supplier context, i.e., SCM, I investigate the empirical case of railroad contracts between Indian Railways and its vendors. I examine how contracts are enforced when a vendor supplied locomotive component fails and some disagreement regarding the cause of failure ensues between the vendor and Indian Railways. Briefly, I find that incomplete understanding of this underpinning complex technical system (Simon, 1962), affects the decision factors (e.g., engineering principles) managers' select. When managers' reason from misplaced decision factors, it leads to flawed conclusions, typically resulting in honest disagreements between the vendor and railroad. The selection of misplaced decision factors is simply a manifestation of boundedly rational managers' attempting to tackle ill-structured problems (e.g., Simon, 1973). Further, in the cases examined in the present study, when both exchange parties have a *shared* but empirically invalid understanding of the underpinning complex technical system, disagreement can be exacerbated and lead to contract termination.

My analysis indicates that contrary to what extant theory suggests (Grover & Malhotra, 2003; Williamson, 1996), opportunism is *not* a necessary condition for contractual disagreements and contract terminations. These findings imply that research has overlooked a separate source of contractual disagreements and contract terminations – uncertainty resulting (unexpected) component failure of the underpinning complex technical system being contracted for effectively. To the best of my knowledge, the current study is amongst the first to examine the implications that has origins in the complex technical system. In light of my findings, I suggest that opportunism is usefully treated as an empirical premise, and contractual parties would do better by not defaulting to invoking opportunism for causes of component failures. In practice, not defaulting to opportunism implies acting in the ‘spirit of the contract,’ rather than enforcing the ‘letter of the contract’ (Williamson, 1996: 115), especially when managers are confronted with ill-structured problems originating in the underpinning complex technical system being contracted for by the exchange parties. I conclude my study by developing refutable anticipated implications regarding the circumstances in which disagreements occur.

2.2. Related Literature

SCM focuses on the “planning and control of materials and information flows as well as logistic activities” (Chen & Paulraj, 2004: 119), for fashioning workable solutions, improving operations (Zipkin, 2012), and facilitating facilitate better coordination between transacting parties (Monczka, Trent, & Handfield, 1998). Although, SCM adopts a system-wide perspective across multiple organizational layers and boundaries (Mentzer, DeWitt, Keebler, Min, Nix, Smith, &

Zacharia, 2001), the buyer-supplier dyad "is of paramount importance to the effective management of supply chains" (Chen & Paulraj, 2004: 120). Specifically, SCM entails a "strategic orientation toward cooperative efforts to synchronize and converge intrafirm and interfirm operational and strategic capabilities" (Mentzer *et al.* 2001: 7). To examine these inter-organizational relationships, research adopts a wide array of theoretical lenses including Transaction Cost Economics (TCE) (Cao & Lumineau, 2015; Villena & Craighead, 2017; Wang, Zhang, Wang, & Sheng, 2016). The use of TCE is hardly surprising – after all, "TCE and SCM are concerned with the same phenomenon" (Zipkin, 2012: 465), and "buyer-supplier relationships, contracting, and internal organization can be effectively examined with a TCE lens" (Anand & Gray, 2017: 5).

2.2.1. SCM and TCE

Research in SCM has employed the TCE lens to highlight the importance of safeguards (Jap & Anderson, 2003) as well as formal and informal mechanisms to curb opportunism (Cao & Lumineau, 2015). This research also shows the existence of opportunism stemming from goal misalignment and other factors affect outsourcing performance (Handley & Benton, 2009), as well as prescribing vertical integration in the presence of uniqueness of design (competence) to guard against opportunism, (Ulrich & Ellison, 2005). Perhaps unexpectedly, it has also been shown that the duration of past exchange relationships does not attenuate opportunism (Mahapatra, Narasimhan, & Barbieri, 2010). In turn, SCM research has enriched TCE by highlighting the salience of bias (Mantel, Tatikonda, & Liao 2006) and factors such as criticality of capabilities in make or buy decisions (McIvor, 2009). The extant literature here also offers evidence of the mediating role of power on exchange relationships (Handley & Benton, 2012) and sheds light on firms' decisions to use intermediaries when entering a foreign market (Kistruck, Morris, Webb, &

Stevens, 2015). Furthermore, it has uncovered a non-linear relationship between international diversification and financial performance (Lampel & Giachetti, 2013).

While the exchange has been mutually beneficial, in SCM, there is a tendency to emphasize opportunism despite the fact that *asset specificity*, not opportunism, is considered to be the "big locomotive" that affects outsourcing decisions (Williamson, 1985: 46). For example, Grover and Malhotra write, "opportunism indicates that human actors in the exchange relationship *will* be guided by considerations of self-interest with guile" (2003: 459, emphasis added). Zipkin notes, "someone *will* be tempted to act selfishly and thus destroy the harmonious accord" (2012: 465, emphasis added). Jap and Anderson suggest, "ex post opportunism will persist in exchange in spite of the firm's best efforts to eliminate it" (2003: 1684). Choi and Krause interpret the fundamental premise underpinning TCE being "industrial players necessarily behave opportunistically" (2006: 644). And, finally, Amundson states, "organizations exist to control opportunism" (1998: 348)

However, TCE does not regard all economic agents as being opportunistic. Rather, it assumes "*some* individuals are opportunistic *some* of the time" (Williamson, 1996: 48, emphasis in original) and maintains opportunism "has reference to exceptions." (Williamson, 2008b: 46). Crucially, TCE suggests that opportunism need not automatically arise because of information asymmetries, nor does it maintain that 'make or buy' decisions are solely determined by considerations of opportunism (Williamson, 1996). TCE is primarily a theory about efficient economic exchange. In the interest of ensuring that *both* parties involved in an economic exchange benefit, TCE advises that transactional attributes (e.g., asset specificity) should be matched with governance structures (e.g., markets and hierarchies) in a manner that economizes on transaction costs. Towards that objective, TCE prescribes incorporating appropriate safeguards because while opportunism may be infrequent, when there is a lot at stake, transacting parties are known to "defect from the spirit of cooperation" (Williamson, 2008b: 46).

Zipkin (2012) notes, some scholars have criticized the excess emphasis on opportunism in OM and general management. Coase suggests preoccupation with opportunism can cause scholars to misinterpret why contracts take the form they do to the extent that "if it is believed that certain contractual arrangements will lead to opportunistic behavior, it is not surprising that economists misinterpret the evidence and find what they expect to find" (2006: 275). Foss and Weber (2016) posit that as a theoretical assumption, opportunism overshadows bounded rationality and reduces the latter to merely an auxiliary assumption and explanation of why contracts are incomplete. Hodgson maintains that in TCE, the role of opportunism has been over-emphasized to the detriment of overlooking "*additional* and very important sources of contracting problems" (2004: 405; emphasis in original). Freeman, Harrison, Hicks, Parmar, and de Colle state that it if the primary concern of (contract) design is opportunism, "one is likely to leave out important ideas such as human dignity, cooperative endeavors, and the creative spirit" (2010:8). The current study will not focus on resolving the debate among organizational economics theorists about whether opportunism is a *necessary* condition for the existence of the firm. However, I seek to demonstrate that managers run the risk of overemphasizing opportunism as the cause of contract failure.

Overemphasis on opportunism by managers has not received much attention in the extant literature on TCE. Indeed as the extensive review of the empirical literature conducted by Crook, Combs, Combs, Ketchen, and Aguinis (2013) illustrates, uncertainty stemming from the behavior of economic agents (managers) is the subject of most empirical investigations. When environmental uncertainty is explored in the extant literature, it tends to take the form of technological uncertainty stemming from high rate of technological change and uncertainty in the timing of the obsolescence of the technology (Balakrishnan & Wernerfelt, 1986; Handley, 2017). However, Alchian and Woodward (1988) suggest that disagreements between transacting parties can occur even when these transacting parties have honest but differing perceptions pertaining to events that occur during an economic exchange. Furthermore, because managers of transacting parties are boundedly rational, “they make mistakes” (Alchian & Woodward, 1988: 66). Congruent with Alchian and Woodward’s (1988) suggestion, the focus of the current study is on ascertaining whether, in practice, contractual disagreements can occur in the absence of transacting parties acting opportunistically, and if so, why, and what the implications are. To do so, as Alchian and Woodward (1988) suggest, considerations of bounded rationality must be enfolded in the study.

2.2.2. Bounded Rationality, Complexity and Uncertainty

While extant research in TCE and SCM does make the distinction between behavioral and environmental uncertainty (Balakrishnan & Wernerfelt, 1986; Grover & Malhotra, 2003), uncertainty stemming from the emergent behavior of the underpinning complex technical system, and its implications for managerial decision making is overlooked. In contrast, the current study explores the possible implications of boundedly rational managers engaged in an economic exchange for complex technical systems. First, managers attributing observed failures of technical systems to the plausible conjecture that such failure was due to opportunism, may be mistaken. Second, managers’ boundedly rationality also affects their treatment of complex technical systems – the standard response being to simplify the system by partitioning it (Dekker, 2011). However, in partitioning complex technical systems, managers may in fact lose valuable information that limits (distorts) their understanding of the system (Dekker, 2011). Due to reasons of bounded rationality, managers' may also be unable to predict the range of failures of a complex technical system, or identify the underpinning causal mechanisms when these outcomes obtain (Holcomb & Hitt, 2007). Managers holding substantially misplaced premises, which can lead to unforeseen outcomes, exacerbate the managerial lack of full understanding of complex technical systems (Augier & March, 2008). The selection of appropriate premises by managers is hindered because ill-structured problems cannot simply be transformed into well-structured problems by the use of heuristics or approximations– and can (often) result in inappropriate decisions (Groop, Ketokivi, Gupta, & Holmström, 2017: Reason, 1997).

To uncover implications (for contracting) of managers grappling with non-behavioral uncertainty originating in the underpinning complex technical system requires an empirical examination of cases where disagreements between transacting parties contracting for complex technical systems occur. However, to identify patterns (of disagreements) and understand why these patterns emerge requires procedural and contextual understanding. To facilitate this 'managers-eye view,' it is necessary to understand the mechanisms used by managers to select their decision factors, while simultaneously incorporating the role of collective rationality (Furubotn & Richter, 2000: 171). In the context of buyer-supplier relationships, some of these decision factors are provided by the choice of governance form (e.g., contract) itself (Grover & Malhotra, 2003). The contract between buyers and suppliers acts as a focal point (Kreps, 1996; Schelling, 1960) for managers involved in economic exchange – a shared premise that encodes their beliefs about how the underpinning complex technical functions, and reflects collective rationality (of the transacting parties) (Simon, 1976); facilities mutual adaptation. While the contract is an instrument designed to facilitate mutual adaptation in the face of uncertainty, it can be poorly designed and misapplied (e.g., design predicated flawed premises pertaining to complex technical systems). The other critical variable used by managers for decision-making is the construct of perceptions about complex technical system failures. In the following section, I discuss briefly how the selection of these decision factors may be influenced.

2.3. Selection of Decision Variable

The selection and application of appropriate decision factors constitutes a managerial challenge when managers' understanding of how complex technical systems work is incomplete (Ramasesh & Browning, 2014; Reason, 1997), and information is not available, especially in case of failures where the component has been extensively damaged. Furthermore, bounded rationality imposes limits on language, which impedes transmission of information, and limits on attention (Williamson, 1975), thereby impeding the managers' ability to reconstruct failures, identify, and apply the appropriate factors in their decision-making. Finally, managers search for information to understand outcomes such as failures may also be affected by their (organizationally induced) biases such as being unable to overcome existing mental schemas (Ramasesh & Browning, 2014), as well as a tendency to conserve cognitive resources (Simon, 1976).

Because managers from differing organizations are likely to possess differing expertise, and be affected by different biases (induced by their respective organizations), there may be a proclivity to use differing decision factors for decision-making. Additionally, the goals of managers of different transacting organizations will differ, which is likely to affect the selection of factors for decision-making to achieve these goals. These differences in decision factors employed by managers of differing organizations are a *regular* feature of economic transactions between parties. These decision factors may converge, or in some cases, diverge, due to myriad reasons such as bounded rationality of agents, differences in knowledge and goals of agents, legalities, ignorance, and opportunism. In addition, contracts are unavoidably incomplete and gaps remain to be filled (Williamson, 1985), and managers across differing organizations may interpret these contracts differently for reasons such as ambiguities in the language of the contract itself (Boardman, 2005).

It is crucial to note that these divergences in decision factors selected by managers of differing organizations in the course of economic exchange do not require any nefarious reasons or ill-intent. Consequently, this study does not assign an *a priori* reason (e.g., opportunism) for divergence of decision factors between managers of different organizations and differing conclusions they (may) reach. The only theoretical assumption employed in this study is that managers are boundedly rational. Adopting this approach allows for the examination of contractual disagreements and contract termination without presupposing human behavioral causes (e.g., opportunism), thus facilitating a deeper understanding of these phenomenon of contractual disagreements and contract terminations.

2.4. Research Methodology

This study uses case research as theory elaboration design in which empirical data are used not only to test managers' causal beliefs but also to challenge them. The approach involves disciplined iteration between causal beliefs and empirical evidence in the context of complex technical system failure, and uses abductive reasoning (Ketokivi & Choi, 2014; Mantere & Ketokivi, 2013). My elaboration of systems failure allows for the combination of several different causal beliefs, which can be generative of new concepts and theoretical insights. Specifically, I use abductive reasoning to draw inferences to explain specific cases and to identify plausible mechanisms that are warranted (Behfar & Okhuysen, 2018; Mantere & Ketokivi, 2013). Identifying these underpinning mechanisms enables generation of a set of boundary conditions for causal explanations (Leibenstein, 1976). Discovering such patterns and causal relationships requires identifying "antecedents" and "consequents" of contractual disagreements following technical systems failure and inquiring which is "connected with which" (Dul, Hak, Goertz, & Voss, 2010: 1175; McLachlin, 1997; Mill, 1882: 470; Van de Ven, 2007). Observation of the outcomes and relationship variables requires some prior theoretical knowledge, and thus existing theory plays a critical role in conducting examinations aimed at producing explanations for technical systems failure.

2.4.1. Background Theory and Sampling Strategy

For the purposes of the current study, TCE (Williamson, 1985) plays the role of background theory. The study focusses on contractual disagreements, contract termination, and the mechanism that leads to them, and posits that there are baseline costs of transacting in conducting routine transactions (Coase, 1937; Williamson, 1975). Contractual disagreements between the railroad and vendor, and contract terminations, are taken as signs of friction between the transacting parties – the critical question the present study seeks to address is - *what causes the 'coefficient of friction' to spike?*

Because variance is essential in identifying causal relations, cases where claims of exchange parties converge and diverge are required (Dul *et al*. 2010; Eisenhardt & Graebner, 2007). The current study focuses on the variance of observed outcomes concerning disagreements between contracting parities, and terminations of contracts. The goal is to identify antecedent conditions for these disagreements and terminations.

2.4.2. Research Setting

This study focusses on buyer-supplier contracts in the railroad industry. The setting is particularly useful because it is through the combination of failure records, official documents, and interviews that the selection of decision factors, limits on cognition of managers, and process of contracting is revealed. Further, this setting is suitable because: (1) the warranty replacement process is part of contract enforcement, and contributes directly to transaction costs; (2) locomotives, especially engines, are paradigmatic complex technical systems with a hierarchical modular architecture (Simon, 1996); (3) there is a high degree of asset specificity and interdependency involved due to the made-to-order nature of manufacturing in the railroad industry, but neither the railroad nor the OEM vendors are fully dependent on each other, i.e., there is no power asymmetry such that one party can coercively prevail over the other; (4) the contracts between Indian Railways and its vendors are predicated on the empirically unestablished premise that locomotive engine subsystems are operationally modular, i.e., failures do not usually propagate from one subsystem to another subsystem; (5) disentangling and identifying the cause of failure is an exercise that makes the limits of individual and collective rationality evident; (6) while contracts incorporate safeguards such as warranty clauses, contractual disagreements, and contract termination occur; and (7) the OEM vendors for engine components have decades of history of transactions with Indian Railways, which informs their contracting and managers' behavior.

2.4.2.1. The Economic System

Indian Railways owns and operates two major locomotive production and assembly plants – Diesel Locomotive Works and Diesel Modernization Works. Together, these production and assembly plants have an annual turnover in excess of USD 500 million, assemble and rebuild over 350 locomotives every year. The majority part of the expenditures of both these plants is the purchase of subsystems and components for assembly of locomotives, and the warranty replacement requests from Diesel Locomotive Works and Diesel Modernization Works to their vendors amount to millions of dollars annually. For safety and quality reasons, government regulations prohibit purchase of components from sources other than vendors that have undergone a rigorous qualification process. Initial approval admits the vendor into a two-tiered system, and order quantities to vendors are allocated based on their tier position.

The two-tiered system works as follows: Tier 1 vendors are collectively awarded 85 percent of the total orders of specific components and subsystems, while Tier 2 suppliers would share the remaining 15 percent. Given the cost of components (e.g., a single turbosupercharger costs in excess of USD 50,000) and volume of production and the facilities, the tier position of the vendor has substantial economic consequences. Ostensibly, the logic of the two-tier system was to ensure high-quality supply while simultaneously attenuating the economic holdup problem. Specifically, qualifying for Tier 1 involved vendors meeting high performance standards over an extended time-period. The economic exchange between all Indian Railways manufacturing and assembly plants and (Tier 1 and Tier 2) vendors is governed by a boilerplate contract, which is modified as required depending on the specific component or subsystem. In the case of a dispute or conflict between the assembly plants and vendors, the matter can be referred to a third party for arbitration, and if the matter remains unresolved, can finally be referred to the Indian judiciary.

2.4.2.2. Operations

Indian Railways owns and operates over fifty diesel depots, which act as maintenance and housing facilities for locomotives. All scheduled and unscheduled maintenance caused by component failure takes place at these depot facilities. The depot staff consists of engineers, technicians, personnel to process and facilitate warranties, logistic specialists and management. In the case of a component failure, the engineers and technicians bring the locomotive to the facility for teardown analysis and repair, and a case is opened in the railroad records. The outcome of the analysis conducted by the engineers constitutes a decision factor from which the manager(s) make decisions (e.g., for railroad managers – whether to request warranty or not). Depending on the severity of the failure, the Indian Railway managers may request an engineer (and manager) from the vendor firm supplying the component to participate in a joint investigation with the railroad engineers to determine the cause. Usually, the Indian Railways managers will take into account the (joint) analysis conducted by the engineers before making the final decision regarding an open case. Cases are considered to be closed once a mutually (between Indian Railways and the vendor) agreed decision regarding the warranty request is reached. When the railroad managers and the engineers from both firms are unable to reach a mutually acceptable technical conclusion for the component failure, this lack of consensus can lead to a contractual disagreement. Indian Railways can then enforce the letter of the contract, and deduct the cost of the failed component from the vendors’ future payments. The vendor can take this action of cost deduction to arbitration.

2.4.3. Data and Preliminary Analysis

Using the Right to Information Act of 2005, I obtained individual component failure records for a fleet of diesel locomotives that were assembled at Diesel Locomotive Works and

Diesel Modernization Works. Component failure data from the former contained approximately 5000 records spanning 2011-2015, while from the latter approximately 1500 records from 2012-2016. These failure data were categorized by major subsystems encompassing several minor subsystems and components. For example, the engine is considered to be a major subsystem, within which power pack constitutes a minor subsystem, and the power pack itself consists of components such as the cylinder liner, piston and cylinder head. For each component failure, Indian Railways records details of components that failed, such as in-service and failure dates, associated locomotive unit numbers, part numbers and serial numbers of components, observations, cause(s) of failures as identified by the railroad engineers, financial data, details of internal and external correspondence, internal remarks for administrative purposes (see Figure 2.1). These records of component failure also include comments for warranty cases where there are delays and disagreements, comments on the status of warranties (such as *"attended," "replaced," "closed," "repaired"*) responses by the OEMs. In cases where the warranty case for the failed component remained unresolved, specific details are recorded.

****

Insert Figure 2.1 here

****

I have several years of experience working in the railroad industry, specifically in teardown analysis of engines, warranty resolution and providing technical input for drafting service contracts. This work experience in the railroad industry was useful in understanding the phenomenon, and interpreting and coding the component failure data. I interacted with Indian Railways engineers and managers to understand the technical and managerial problems they encountered in assembling and operating the locomotives, and continued email correspondence as well as telephone conversations with them after my visits. I also interacted with the experts in the warranty department at both plants to ensure proper interpretation of the data. These interactions also helped me understand the internal (to Indian Railways) and administrative aspects of record keeping, warranty-resolution, and procure for closing warranty cases. I was also able later to seek the guidance of the warranty department experts in coding data.

In addition to failure data, I also obtained a copy of the agenda of the annual (4th) quality meeting held by the Indian Railway's Diesel Maintenance Group in 2010. This meeting pertains to all abnormal failures and persistent quality issues faced by maintenance depots across India. The issues are subdivided by major components and subsystems such as engines, traction motors, bogies, and each issue is summarized, specific problems are highlighted, actions taken are recorded with cross-references to other meetings held at various assembly plants and depots across India. This collated record was critical in identifying major and persistent quality problems and failures.

2.4.3.1. Preliminary Analysis

My initial analysis revealed that at the time of data collection, approximately 13 percent of the all warranty replacement cases involved some form of disagreement (but no arbitration) between Indian Railways and its vendors. The substantial number of cases of disagreements between the transacting parties indicates that there was friction in enforcement of the contract. By cross-referencing failure records and the Diesel Groups meeting agenda, I was able to code the data in a manner such that persistent failures and disagreements were brought into focus. The primary purpose of this coding was to narrow down the inquiry to delve deeper into specific outcomes of interest to better understand the process. This study primarily focusses on cases where there were persistent failures and some (technical) disagreements and contract termination. Such disagreements indicate an incomprehensive understanding of the underpinning complex technical system. In addition, these disagreements between Indian Railways and its vendors also indicate the presence of (some) failures patterns that were part of recurring transactions pertaining to particular components. Despite recurring component failures, transactions between Indian Railways and vendors supplying the failed component typically persist due to the mutual dependency between both transacting parties arising from high relationship-specific investments (Williamson, 1975, 1985). The current study specifically focuses on mechanical failures because the technology associated with mechanical subsystems is relatively stable. Moreover, the failure data on electrical subsystems tends to be aggregated at higher levels of system architecture and often lacks granularity thereby often impeding an understanding of the causes of failure. Within mechanical systems, engine related component failures featured prominently.

I then requested, and was given access to, several internal documents relevant to engine component related issues including a copy of a presentation given by the OEM management and engineers to Diesel Locomotive Works experts in 2009. I matched these internal documents to the failure records, and was able to identify specific components where disagreements and discord were common. These disagreements about the cause(s) of component failure were evident by remarks recorded by Indian Railways for internal administrative purposes. To ensure that I had sufficient technical understanding of the underpinning technical system, i.e., the locomotive, I obtained the engineering drawings of components, as well as assembly manuals and maintenance instructions provided by one of the OEM vendors. These OEM manuals and instructions provided by Indian Railways were very useful in understanding the reason for disagreements pertaining to expected standards of assembly, performance, and maintenance. Crucially, I also obtained the contract signed by Indian Railways and its vendors, which covers warranty clauses, procedures to be followed by the railroad and vendors to process and close warranty requests, and conditions for invoking arbitration and settlements.

2.4.3.2 Interviews

In 2015 and 2016, I visited both Diesel Locomotive Works and Diesel Modernization Works. During my visit to Diesel Locomotive Works, I interviewed the head of the production department as well as the chief engineers at the engine, turbocharger design, and production divisions. Both interviews were semi-structured and lasted over two hours each. The intent of the interviews during the course of the visits was to elicit detailed responses of cases where contractual disagreements or conflicts had occurred, and how these were ultimately settled regardless of whether the cases remained resolved or not. In the course of these interviews, I was provided an internal presentation given by experts at Diesel Locomotive Works to the Railway Board in 2015. Throughout the visit, I made field notes. To mitigate any misunderstandings, after my visit I followed up via email correspondence with the production head of the facility.

I visited Diesel Modernization Works twice, and on each occasion, I interviewed the CEO. For both interviews, the top technical experts and design engineer of the plant were invited to participate. I followed the same pattern of inquiry during the visits, and made notes. I followed up these visits with a telephone interview of the CEO and written email correspondence with the top engineers. In my second visit, I was given access to correspondence, in the form of letters and faxes, between the Indian Railways and OEM where a disagreement arose. This correspondence record included engineers and managers comments from both organizations as they used technical and legal (contractual) arguments to convince the other party of their conclusion(s).

****
Insert Figure 2.2 here
****

During the course of my visits, it became evident that managers at these locomotive production and assembly plants have a strong grasp of contracting-related problems. For example, without using academic language, the managers expressed concerns about opportunism hazards. The institutional and political environment also explicitly prescribes that managers mitigate such potential contractual hazards (e.g., the initiative to find and develop local vendors to ensure that foreign vendors do not “hold-up” production). Indian Railways management is also quite familiar with engineering related issues, including technical details, due to railroads’ policy of promoting engineers through the ranks.

Since even very senior managers at the Indian Railways assembly and manufacturing units may not have a complete grasp of the (strategic) goals of the railroad as an organization, and their interpretation of the contract may differ from the governing board, I met with the Additional Railway Board Member (retired) and Full Board Member (retired). The latter was the sole in charge of all Indian Railways activities related to diesel locomotives, while the former was also in the same division, and hence intimately familiar with the setting. These meetings with the retired board members took place on separate occasions, and the semi-structured interviews lasted approximately three hours each. I took notes at the interview, and followed up by email correspondence. Because these interviews were conducted separately, I was able to corroborate their responses and attenuate concerns of recall biases.

As my data collection involved several rounds, I was able to identity and focus on cases where there was agreement or disagreement between the contracting parties, or a contract termination. This strategy allowed me to delve deeper in the data such that scenarios emerged. By scenarios I mean a pattern of cases that were similar in technical failures and outcomes, and could be grouped together. This approach of grouping lead to aggregation and allowed the isolation of exemplar cases illustrating the scenarios of agreement, disagreement between transacting parties, and contract termination. In the following section, I discuss contextual factors pertaining to the setting, and then proceed to a discussion of the scenarios.

2.5. Case Analysis

At the outset, it is important to highlight certain contextual factors since these may affect the behavior of actors of organizations and any subsequent inferences drawn. These contextual factors that need to be borne in mind are as follows. First, while typical vendors supply components for only one type of locomotive, Indian Railways typically has multiple models of locomotives being maintained at the same depot with the same personnel responsible for their operation and upkeep. Second, there are safety considerations involved which impede engineers from investigating certain failures (e.g., replicating some failures may be dangerous). Third, Indian Railways is a government organization and subject to the general by-laws of the government of India, and is driven by its own internalities and broader political goals (Wolf, 1993). These goals can make Indian Railways seem rigid, but certain procedures have been designed as safeguards, even at the expense of efficiency. It does not follow that there is no room for improvement within these procedural constraints. Finally, there is a boilerplate contract between Indian Railways and its vendors that stipulates supplied components are covered by warranty for *"165,000 kilometers after commissioning of the locomotive or Twenty Four (24) months after receipt"* (commissioning refers to the date the locomotive was put in service).

In a typical case of a vendor supplied component failure, the depot opens a case and records the locomotive number, date of failure, suspected cause and details, serial number, part number etcetera and sends the report and warranty replacement request to the vendor. Upon receiving these details, the vendor can accept the warranty request, ask for an inspection or additional data, or, deny it. The acceptance of a warranty request on receipt is treated as the routine (baseline) case of contractual agreement. Cases of contractual agreement between the transacting parties include instances where some additional data is requested by either organization, or some technical clarifications are required. Cases where transacting parties are unable to converge to (technical) conclusions despite the provision of additional data, and additional resources are consumed in attempts to coalesce to some mutual understanding are considered to be contractual disagreements. Such disagreements tend to be recorded and addressed in venues such as Indian Railway's Diesel Maintenance Group meetings, are escalated to the Railway Board level, and depart from the agreed upon period for resolution of warranties (usually 120 days). According to Indian Railway experts, the opportunity cost of delays in availability of locomotives is estimated to be roughly USD 6,000 per day, while the current production staff to administrative staff ratio in the production plants is almost 1:1, a figure that is certainly higher than similar railroads elsewhere, and considered to be excessive even by top Indian Railways management. These figures can be considered indicators of transaction costs, in that they reflect the underpinning cost of exchange (Coase, 1994: 7-8; also see Merkert, Smith, and Nash, 2012 for details on administrative costs representing transaction costs in the context of railroads). I now proceed to discuss the five scenarios, of which, one involves contractual agreement and serves as a baseline case, and four involve either contractual disagreements or contract termination.

2.5.1. Scenario 1: Agreement between transacting parties

This scenario of contractual agreement between Indian Railways and its vendors pertains to cases where the warranty processing is straightforward, and the transacting parties converge to a conclusion pertaining to failure (e.g., vendor grants warranty request for a failed component thereby accepting a quality problem). For example, warrant case records obtained at Diesel Modernization Works indicate that on a locomotive commissioned on 22nd February 2014, the reverser (gear lever for a locomotive used to control its movement) failed on 29th December 2014. The warranty request was registered on 12th January 2015, and the warranty case report noted that the technical problem was *"contact burnt."* The vendor replaced the defective component on 12th February 2015. Both Indian Railway and the vendor of the failed component agreed that there was likely a quality problem associated with the supplied component that lead to its failure. In another case of component failure, a turbocharger compressor was found *"damaged."* The date of turbo-charger compressor failure recorded by Indian Railways was 7th October 2011, while the commissioning date was 12th September 2009. Indian Railways registered the warranty request with the vendor on 9th June 2012, and the vendor denied the request stating, *"warranty is not tenable due to service life is more than 24 months."* The warranty case was closed on 13th August 2012. In this particular case, the vendor drew Indian Railways' attention to the fact that the failed component was not covered by the contract, and Indian Railways accepted the denial of their warranty request and closed the case. In two other separate cases of warranty replacement requests by Indian Railways to a vendor for Cam Rollers with *"hardness less"* being cited as a cause, an investigation was conducted, and the warranty request was denied by the vendor because a recheck reading showed the Cam Rollers were *"ok as per specification."* Indian Railways accepted these readings provided by the vendor, and the warranty cases were closed.

In all these warranty cases, the premises were well established, and routine procedure was followed by the transacting parties. At various times in the contract enforcement process, observation of failure, engineering principles, and the contract were used as decision factors, and the relevance and applicability of these factors was unambiguously accepted by the transacting parties. Managers of these transacting parties were dealing with well-structured problems (Simon, 1973), and contract enforcement was straightforward.

2.5.2. Scenario 2: Contractual disagreement - case of turbosupercharger assembly

This scenario of contractual disagreement refers to a disagreement between Indian Railways and an OEM vendor that occurred in 2009 concerning specifications for turbo-supercharger assembly. The assembly specification that lead to the disagreement between the Indian Railways and the OEM vendor was the impeller eye (rotating part, like a fan) and casing clearance (gap) at the six o'clock position. The allowed gap specification is between five and twenty-five thousandth of an inch, which is less than the thickness of a human hair. These gaps in turbosupercharger assemblies are so tight, that during assembly, even one extra turn of the screw can lead to the turbosupercharger not meeting the allowed specification. In fact, to ensure that contaminants do not affect assembly tolerances and the required assembly specifications are met, Indian Railways assembles turbosuperchargers in positive pressure rooms that expel contaminants via the ventilation system. Records from a meeting held between Indian Railways and the OEM vendor in March, 2010, indicate that the railroad was using the OEM vendors maintenance instructions for assembly, which are laxer in allowed tolerances and easier to follow than instructions issued for assembly. At the time of the meeting, poor performance had caused Indian Railways to overhaul prematurely 29 out of the 371 turbosuperchargers they had assembled. The same meeting records also indicate that the OEM vendors' top management and senior engineers' highlighted deviation from the assembly instructions would lead to *"lowered turbocharger efficiency on new turbos."*

However, Indian Railways thought it was perfectly acceptable practice to use maintenance specification for assembly purposes, and their primary counter argument against the OEM vendor being that maintenance instructions were applicable immediately after the first post-assembly field use of locomotives. Citing engineering drawings of the turbosuperchargers, Indian Railway managers also claimed that the OEM vendor had neglected to provide "*maximum or minimum clearance for this inspection…meant to verify proper turbo installation not proper impeller eye clearance.*" The OEM vendor managers insisted that initial assembly conditions affect performance and failure in the long-run. While they admitted that the assembly instructions did not contain a separate installation verification procedure, the OEM vendors' managers noted that the instructions issued by the OEM did contain precise assembly specifications. These specifications provided by the OEM vendor included final tolerances, and the problems faced by Indian Railways were likely a result of improper assembly and not meeting those stipulated standards. The transacting parties eventually negotiated a solution that involved the OEM vendor conducting a hands-on training and audit of the process followed at the Indian Railways facility at the vendors' cost. After these training sessions, Indian Railways agreed to maintain, during assembly, the clearances, and monitor additional parameters, as prescribed by the OEM vendor.

In this particular disagreement between Indian Railways and the OEM vendor, it was arduous for the OEM vendor to address directly Indian Railways' question regarding why maintenance instructions ought to be followed as soon as the locomotive leaves the manufacturing facility, but not in the assembly phase itself. Even assuming the OEM vendor could address the question, clearly establishing the implications of deviating from assembly instructions would have been a considerable engineering exercise involving substantial expenditure of resources. Further, certain knowledge of assembly procedures cannot be transmitted except by doing (Polanyi, 1958). In this case of disagreement over the assembly procedures, not only were managers' dealing with an *ill-structured problem* (Simon, 1973) in that premises and conclusions pertaining to variables affecting performance of the turbosupercharger could not be unambiguously established, there were limitations of language as well. During the semi-structured interviews, Indian Railways managers revealed that such disagreements between the railroad and OEM vendors were usually resolved after considerable dialogue, and negotiation training where required. The provision of training by the vendors' is typically *not* contractually stipulated. This case serves as an example where the contractual disagreement was finally resolved by mutually agreeing to extraneous (to the contract) measures, and not by definitively addressing the technical problem originating in the underpinning complex technical system.

2.5.3. Scenario 3: Contractual disagreement - case of premature piston ring failures

This scenario of contractual disagreement refers to a disagreement between Indian Railways and an OEM vendor over the premature failure of piston rings. Records from an internal Indian Railways quality meeting held in 2010 indicate that the railroad managers were concerned with early failures of piston rings. The components were failing after two years, which though beyond the contractual replacement terms of two years, were a source of persisting tension between the OEM vendor and Indian Railways because engine component life is rated to be six years. Notes from a meeting and presentation between Indian Railways and the OEM vendor held in May 2010 records the OEM vendor response being *"[they were] struggling to understand this concern"* and *"similar vintage piston rings on other accounts"* were *"getting proper life."* The OEM vendor requested additional data from Indian Railways concerning each parameter outlined in the maintenance manuals, and were specifically concerned with the oil consumption of the locomotive engine and the associated filter changes.

An interim investigation conducted by Indian Railway engineers on early wear of the piston ring lead them to hypothesize that the problem of these early failures may be caused by inadequate hardness of the piston ring. The Indian Railway engineers subsequently followed that line of investigation to confirm their hypothesis. In contrast to Indian Railways engineers, guided by prior experience, the OEM vendor engineers checked spectrographic reports of oil circulating around the piston ring, and narrowed down the plausible cause to excess silicates (sand) in the circulating oil. Both these differing explanations regarding the cause of piston ring failures developed by the Indian Railway engineers and the OEM vendor engineers were plausible, and not possible to dismiss. In fact, premature piston ring failures could plausibly be explained by the conjunction of the two factors. Managers of the Indian Railways and OEM vendor managers were faced with an *ill-structured problem* (Simon, 1973). In this case of premature failure of piston rings, the premises applied by the Indian Railway managers and the OEM vendor managers, because of their differing expertise were different, subsequently leading to differing conclusions. Eventually, continuing investigations conducted by the parties in which data was shared revealed that Indian Railways maintenance personnel were using instructions for older model locomotives. These instructions for older locomotives allowed a much higher concentration of contaminants such as silicates in the oil than the specifications for the OEM vendors' newer supplied engines allowed, and the presence of these contaminants (in excess) could plausibly explain premature piston ring wear. The contractual disagreement serves as an example where a disagreement was finally resolved by identifying a deviation from procedure (stipulated maintenance practices) by an organization (Indian Railways), and not by definitively addressing the technical problem originating in the underpinning complex technical system in that the causal link between this deviation and premature failures was not empirically established.

2.5.4. Scenario 4: Contractual disagreements – case of locomotive engine quality

This scenario of contractual disagreement pertains to a disagreement between Indian Railways and its vendor over of lube-oil-to-fuel-oil consumption ratio and locomotive engine quality. The lube-oil-to-fuel-oil consumption ratio is considered to an excellent indicator of the 'health' of high-horsepower locomotive engines, and deviations on the higher side are a cause of concern. This consumption ratio is encoded into the performance standards stipulated in the contract between the Indian Railways and vendor, and is usually tested by the railroad. In a contractual disagreement, which occurred between the railroad and a Tier 1 OEM vendor in 2010, Indian Railways claimed that the OEM vendor supplied high horse-power engines were sub-standard as they were not meeting the contractually stipulated performance parameters in the pre-installation (on the locomotive) test in the engine test shop at the assembly plant at Diesel Locomotive Works. This failure to meet the contractually stipulated performance, Indian Railways maintained, led to excessive operational costs, including possibly premature turbosupercharger failures and replacements of power assemblies beyond the warranty period of two years, but earlier than the six years rated life of engine components. Indian Railways shared the performance and failure data with the OEM vendor, so both railroad and OEM were operating from the same data.

However, the Tier 1 OEM vendor denied that there was a quality problem with the engines, and maintained that the engine test shop at their own production facility used to develop the standards was a highly controlled environment, while Indian Railways had an open test shop within the assembly plant. Hence failure to replicate the Tier 1 OEM vendor engine test shop performance by the railroad was not a valid indicator of engine quality. The vendor suggested that because the technical operating specification for the mix of lube oil and fuel oil were not met, substandard performance and premature failures were the result of poor maintenance practices, and highlighted that the same engines were supplied globally without causing similar problems. Further, they stated that the standard stipulated in the assembly documents was a metric for the engine test shop, not locomotives in field operation. Hence, the vendor reasoned, failure to replicate the Tier 1 OEM vendor engine test shop performance by the railroad was not a valid indicator of engine quality and could not be linked with field performance. The OEM vendors' stance was that there was no problem with the engine quality, early engine related failures could be attributed to factors peculiar to India, including possibly poor maintenance practices followed by Indian Railways.

Indian Railways then threatened to downgrade the Tier 1 OEM vendor from Tier 1 to Tier 2, and since there was an alternate qualified OEM vendor, which could be promoted, the threat was credible. Further, Indian Railways produced performance reports of engines supplied by the Tier 2 OEM vendor that claimed to outperform the Tier 1 OEM vendor builds. However, there was a critical factor in the data supplied by Indian Railways, which was highlighted by the Tier 1

OEM vendor – the Tier 2 OEM vendor supplied engines with ‘superior’ performance were newer than the Tier 1 OEM vendor supplied engines, and hence the comparison was not valid. A joint investigation involving Indian Railways engineers and the Tier 1 OEM vendor engineers ensued, which included identifying ten locomotives of similar vintage and mileage profile for the purpose of performance comparison between the Tier 1 OEM and Tier 2 OEM vendors’ supplied engines. Indian Railways and the Tier 1 OEM vendor decided that for the comparison between Tier 1 OEM and Tier 2 OEM vendor supplied engines, the Tier 1 OEM would supply components that were supposed to be regularly changed out (such as filters) and deploy its own engineers to monitor the locomotives, and the Tier 2 OEM vendor could also follow the same practice. The OEM vendor(s) would also be given access to any requested data so that they could monitor factors that affect the performance of the engine.

Subsequently, the Tier 1 OEM vendors’ engineers identified metrics that did not meet their stipulated operating standards (specifically, elevated levels of contaminants in the lube oil), which allowed the vendor to claim that the deviations in lube-oil-to-fuel-oil consumption ratio were due to poor maintenance practices. Indian Railways counterclaimed that those identified factors could not have a severe impact, and that the engines supplied by the Tier 2 OEM vendor were also subject to similar conditions and practices. These investigations of lube-oil-to-fuel-oil consumption by engines lasted over two years and the matter not resolved on the basis of claims based on engineering premises. For each claim advanced by either transacting party, the other to cast doubt on it advanced a counterclaim. Both Indian Railways and the Tier I OEM vendor were unable to establish the veracity of their respective claims that they advanced because the identified metric, the lube-oil-to-fuel-oil consumption ratio is meant to serve as an indicator for the health of the locomotive engine. Given the complexity of a locomotive engine and the operational environment, it was arduous managers to definitively establish a causal relationship, or lack thereof, between a failure to meet contractually specified lube-oil-to-fuel-oil consumption ratio in the test shop and actual lowered field performance and early field failures. In essence, managers from Indian Railways and the Tier I OEM vendor were dealing with an *ill-structured problem* (Simon, 1973).

Ultimately, this matter of lube-oil-to-fuel-oil consumption ratio was escalated to the Railway Board level, and in response the Tier 1 OEM vendor realizing the gravity of the situation, deployed one of its top engine experts to India. The disagreement between the Indian Railways and the Tier 1 OEM vendor was finally resolved when the Railway Board introduced a new policy requiring the railroad to shift from older model lower locomotives to high horse-power diesel engine locomotives due to the latter’s superior hauling capacity. The reasoning, the retired Railway Board members (who had made the decision) revealed in the interviews was that higher horse-power capacity frees up sections of track for use, thereby leading to higher revenues. The economic concern about excess lube-oil-to-fuel-oil consumption ratio (as long as it did not definitively and directly lead to early field failures), the retired Board members revealed, was secondary since in light of the policy shift, the high horse-power locomotives were a scarcity. Thereafter, the excess lube-oil-to-fuel-oil consumption ratio became a maintenance, and not performance, issue. The change in policy regarding the salience and monitoring of lube-oil-to-fuel-oil consumption ratio serves an example where an ill-structured problem that lead to a contractual disagreement between transacting parties was resolved by a change in the organizational policy introduced at a higher level of authority that converted a technical problem associated with the underpinning complex technical system into an issue adequately addressed by routine (maintenance).

2.5.5. Scenario 5: Contract termination – case of engine liner failures

This scenario of contract termination warrants refers to the case cited in the introduction where even after the vendors were swapped, unexpected premature power assembly failures continued to occur. During my visit to Diesel Locomotive Works, the Chief Design Engineer at the Engine Shop brought up issues related to warranty claims for power assemblies, specifically liners. For liners, due to the high asset specificity required for manufacturing, at that time there were only two OEM vendors, one in each tier of the two-tiered system. In 2012, a maintenance depot reported that the failure rate for liners supplied by the Tier 1 vendor was 72%. Concerned about poor quality of supply, the railroad downgraded the Tier 1 vendor to Tier 2, and upgraded the Tier 2 vendor to Tier 1. For this scenario, the downgraded vendor is referred to as Vendor A, and the upgraded vendor as Vendor B. Prior to the spate of failures, Vendor A supplied 85% of the liners, while Vendor B supplied the remaining 15% of the total supply, to Indian Railways.

Examination of data revealed a troubling feature: an empirically unestablished premise was codified as one of the central clauses into the contract between Indian Railways and its vendors. This central clause in the contract states: "*Contractor will not be liable for damage to its equipment consequent to failure of any other equipment not supplied by the contractor.*" This clause reflects an underpinning premise of both transacting parties – that locomotive engines are operationally modular, i.e., failures occur independently of each other – and mutually accepted by the transacting parties involved. Both transacting parties expect that consequential damages from failed components can occur occasionally, but since failures are independent of each other, not systematically so. This belief of operational modularity of locomotive engines was corroborated in my discussions at both the manufacturing facilities, and the fact that this is a standard industry clause. It is important to mention that both contractual parties also anticipate that the clause is in (both) their economic interest: the vendors expect that the clause protects them from being held liable for consequential damages, while Indian Railways (railroads) anticipate that it prevents vendors from charging excessive prices.

In the failure analysis of liners supplied by Vendor A, using the premise of operational modularity, the railroad engineers calculated the failure ratio of liners based on failure count and vendor supply. The engineers' presentation of the analysis to the mangers reflected these finding where failure rates were calculated as if the liners functioned in isolation and were fully operationally modular. The decision to swap vendor tiers was made based on this analysis by Indian Railways engineers. Subsequently, liners supplied by the vendor upgraded to Tier 1, i.e., Vendor B, exhibited an even higher rate of failure, which was attributed to *"variable thickness"* of the liner wall – a quality defect. Indian Railways conducted a failure analysis of these liners supplied by Vendor B, and in doing so, employed the premise of operational modularity again. Indian Railways also recorded that *"corrective action"* taken by the depot that had initially reported high failures railroad included the increased frequency of filter replacement, which had *"significantly"* reduced failures of liners supplied by vendor A. In 2014, the Indian Railways swapped tiers and supply contracts between the two vendors, restoring the pre-2012 status quo. An extensive teardown analysis of failures of liners supplied by Vendor B had led Indian Railways to conclude that the cause of failure was excessive heating of the liner, which *"breaks the lubrication between liner and piston."* Given the extremely tight tolerances involved (less than one-tenth of an inch), excess heating of the liner due to inadequate lubrication caused by poor filters can also explain why variable wall thickness obtained.

Indian Railways did recognize, in passing, that the cause for excessive heating of the liner may have been systematic interactions between components (filters and liners), which is inconsistent with the language of the contract and thus the warrant used for their analysis. Engineering and legal issues aside, this is a case where a false premise – that engines are operationally fully decomposable – was accepted by Indian Railways (and the vendor) led to contract termination. The thickness of the liner wall, which has a critical tolerance lower than one-hundredth of an inch, can vary due to a manufacturing defect, erosion from wear and tear occurring over time, or both. Indian Railways attributed the variable thickness of the liner wall to poor quality of manufacturing, the vendor denied any quality issues. Indian Railways then, without any evidence, *assumed l*ack of candor on part of the vendor, which contributed to ensuing disagreement being escalated and resulting in contract termination. This contract termination occurred because Indian Railways used opportunism as a premise to take action against a vendor they perceived as acting dishonestly, and enforced the contract strictly. The course of action followed by Indian Railways led to substantial ex-post transaction costs for *both* transacting parties. Yet, neither transacting party had actually engaged in any opportunistic behavior – the disagreement between Indian Railways and the Tier 1 OEM vendor was rooted in being unable to satisfactorily explain premature failures of liners, which, in isolation, are generally well understood.

As Indian Railways noted, the technical fault was finally traced to improper filter change-outs, which affects the quality of oil, and therefore erosion time. The initial conclusion Indian Railways arrived at was flawed because they reasoned from a false premise of operational modularity of locomotive engines (Bhardwaj & Ketokivi, 2016). The bounds on individual and collective rationality are evident, manifested in the lack of attention to pertinent information (e.g., increased frequency of filter replacements), selection, and use of convenient unfounded premises (e.g., operational modularity, opportunism) to arrive at conclusions by Indian Railways managers. The case of liner failures is an example of honest disagreements and contract termination when an undesirable unexpected outcome was misinterpreted by Indian Railways (OEM vendor) due to reasons of managerial bounded rationality and incomprehensive understanding of the complex technical systems. The cause of contract termination was reasoning from a mutually accepted empirically invalid premise by both Indian Railways and the vendor, Indian Railways adducing the empirically invalid premise of opportunism to the vendor, and applying the contract strictly. This is an example where Indian Railways managers were dealing with an *ill-structured problem* (Simon, 1973), and the action of contract termination neither addressed the technical problem associated with the underpinning complex technical system, nor did it resolve the contractual disagreement. However, it should be noted that Indian Railways restoring status quo can be interpreted as *mea culpa*, and indicates that opportunism was not at play here.

2.6. Discussion

Careful consideration of the five scenarios discussed in the study reveals a critical pattern. First, the scenarios can be grouped together based on whether managers are dealing with well-structured or ill-structured problems when trying to establish the causes of failure or inadequate performance of the components. In the case of well-structured problems, decision factors and conclusions for failures are clearly identifiable, and there is no ambiguity associated with selecting and establishing them. In the latter case of ill-structured problems, selecting decision factors for the reasons of failures constitutes a non-trivial obstacle not easily overcome, and leads boundedly rational managers to conclusions that cannot be unambiguously established on technical grounds.

****
Insert Table 2.1 here
****

Second, across organizations, decision factors can either be compatible or incompatible, and incompatibility arises as a result specialized knowledge, limits of communication and language, adduced assumptions, and other factors. The recognition these decision factors may differ across organization for a host of reasons does not deny the possibility of opportunism, but merely recognizes that information asymmetry does not automatically lead to opportunistic behavior. Third, the likelihood of disagreements, i.e., friction between transacting parties is affected by whether transacting parties agree or disagree about decision factors *and* by the empirical validity of premises employed in reasoning about failures. In the latter case, where the empirical premise regarding the underpinning complex technical system the parties are contracting for is false, it leaves room to invoke additional premises regarding the *behavior* of the transacting parties. For example, flummoxed by the unexpected spate of failures of engine liners, Indian Railways managers assumed that the OEM vendor had acted opportunistically, despite no evidence of opportunistic behavior. Invoking this premise allowed managers to reduce their search cost for technical reasons for failure, and conserved resources (e.g., cognitive resources), but turned out to be a myopic adaptive response.

An analysis of the five scenarios suggests that depending on whether managers are confronted with well-structured or ill-structured problems, differential likelihood of contractual agreements and disagreements between transacting parties, and contract termination can be expected. Disagreements and contract termination occur because the cause of failure of the underpinning complex technical system, i.e., locomotive engine, cannot be established definitively.

Disagreements do not yield technical solutions, and their remedies are either found in adaptations centered around the contract (Scenarios 2,3, and 4), or can lead to contract termination (Scenario 5). In addition, whether managers act in the ‘spirit of the contract’ or adhere to the ‘letter of the contract’ when they are dealing with these ill-structured problems originating in the underpinning complex technical system being contracted for also affects contractual outcomes. In the following section, the implications of these managerial choices are discussed.

2.6.1. Likelihood of contractual agreements

Contractual agreements represents routine transactions involving baseline transaction costs of contracting and minimal friction, and are associated with managers confronting well-structured problems. Scenario 1 falls into the category contractual agreements. When transacting organizations advance claims and counterclaims involving well-structured problems originating in the underpinning complex technical system being contracted for, compatible decision factors and compatible conclusions are readily established, contractual enforcement is routine, and no disagreements ensue. Obtaining contractual agreements is not always trivial and requires that additional decision factors (e.g., additional data) that are mutually acceptable to both transacting parties be admitted. Further, because transacting parties are dealing with well-structured problems, the claims put forward by transacting parties are likely to be empirically valid, managers can enforce the ‘letter of the contract’ without contractual problems, which is observed in the timely resolution of these warranty claims advanced by the buyer. Therefore, I propose:

> *Proposition 1: An increase in enforcement of interfirm contracts by the letter of the contract between managers confronted with well-structured problems of technical failure within an underpinning complex technical system, decreases the likelihood of interfirm contractual disagreements (over warranty claims).*

2.6.2. Likelihood of contractual disagreements

Contractual disagreements represent transactions involving friction, and are associated with managers confronting ill-structured problems. Scenarios 2, 3, and 4 fall into the category of contractual disagreements. When contracting parties' advance claims and counter-claims associated with ill-structured problems originating in the underpinning complex technical system being contracted for, incompatible decision factors and incompatible conclusions obtain. In each of these scenarios, the complexity of the underpinning technical system is prominently impactful.

Consider scenario 2 pertaining to the case of turbosupercharger assembly and premature failures (overhauls), because the turbocharger is a complex technical system, the decision factors applied by Indian Railways managers to explain these outcomes of premature failures is ambiguous. Yet, the managers of the OEM vendor are unable to dismiss the use of such ambiguous decision factors because it is arduous, if not impossible to show empirically the application of these factors is flawed. Simultaneously, the managers of the OEM vendor cannot established definitively, the validity of their own decision factors. This contractual disagreement pertaining to premature failures of turbosuperchargers was resolved by negotiating in the 'spirit of the contract,' and a joint approach (of training) to address the technical problem.

In scenario 3 pertaining to premature piston ring failures respectively, Indian Railways and the OEMs used differing decision factors because of managers (engineers) differing expertise and competence rooted in their respective organizations, as well as inadequate comprehension of the underpinning complex technical system, i.e., the locomotive engine. Indian Railways' past organizational experience with the technology of the piston ring also played a role in further diverging the managers choice of decision factors. Because of the differing expertise of Indian

Railways' managers and the OEM vendors' managers, and their respective (different) limits of understanding of the complex technical system, different conclusions reached, which lead to disagreement. The disagreement was resolved when it was shown that Indian Railways technicians were deviating from the expected routine (maintenance practice). This identification of deviation from the routine did not necessarily yield a technical solution because it could not be definitively established that the deviation was the sole cause of premature piston ring failures. The resolution to the disagreement was obtained by transacting parties pooling their data, findings, and expertise in the 'spirit of the contract.'

In scenario 4 pertaining to the case of locomotive engine quality. The multiplicity of possible explanations, presence of feedback loops to the system, temporal entanglement of events, and inability to separate relevant variables, impeded the establishment of a definitive explanation for the out of specification ratio lube-oil-to-fuel-oil ratio, and its perpetrated link with locomotive engine performance (quality). This disagreement did not yield a technical solution, but was resolved when the organizational goal of transporting higher volume of freight took precedence over suspect engine performance, and in the 'spirit of the contract,' the (supposed) performance related issue was transformed into a maintenance issue.

These cases share several commonalities. First, ill-structured problems originating in complex technical systems tend to lead to disagreements between transacting parties. Second, these problems tend not to yield definitive technical solutions. Third, in each of the cases discussed, contractual disagreements are resolved via managerial, and not technical solutions, involving managers acting in the 'spirit of the contract.' Such resolution is achieved by, as an executive cited by Macaulay noted, "if you keep the lawyers and the accountants out of it" (1963: 61). Hence,

> *Proposition 2: An increase in enforcement of interfirm contracts by the letter of the contract between managers confronted with ill-structured problems of technical failure within an underpinning complex technical system, increases the likelihood of interfirm contractual disagreements (over warranty claims).*

2.6.3. Likelihood of contract termination

Contract termination represents transactions involving unsurmountable friction between transacting parties, and is associated with managers confronting ill-structured problems. Scenario 5 falls into the category of contract termination. In scenario 5, pertaining to premature liner failure case, Indian Railways managers use premise (shared by the OEM vendor, and encoded into the contract) that the locomotive engine is an operationally modular system while conducting failure analysis of liners. The practical ramification of the assumption of operational modularity is that because failure involving systematic propagation of failure from one component is assumed not to occur, failure analysis is conducted by isolating the component and ignoring any possibility of interactions with other components. The findings of the failure analysis, predicated on the empirically flawed assumption of operational modularity, concluded that the liner wall was a cause for concern. Indeed, the liner wall was found to be damaged due to excessive heating. While Indian Railways concluded that the damage to the liner wall was due to poor supply by the Tier 1 OEM vendor who had been extensively vetted, the vendor vehemently denied any quality issues. Indian Railways managers', without supporting evidence, assumed that the OEM vendor had acted opportunistically, and downgraded it to Tier 1, terminating the contract. In an attempt to resolve the assumed problem associated with liner supply quality, Indian Railway managers upgraded the Tier 2 vendor to Tier 1. Problems with premature failures of liners persisted with the upgraded vendor, and Indian Railway management eventually traced these problems to the filter.

Prima facie, the disagreement between Indian Railways and the downgraded vendor suspected of opportunism seems to be a manifestation of bounded rationality. First, partitioning complex technical systems for analysis is suggested precisely due to limits on human cognition (Simon, 1996). Second, not engaging in a full search for the technical cause of the problem also indicates that bounded rationality was at play. Third, evidence that the premature liner failures were not due to quality of supply was overlooked, likely because limits of attention did not allow managers to register or link the remedial action of frequent filter changeouts with drop in liner failures. The eventual restoration of the downgrade vendor to Tier 1 status suggests "frailty of motive and reason" (Simon, 1985: 303). Contract termination occurred because boundedly rational managers ascribed opportunism to their transacting partner, without any evidence of such behavior, and chose to enforce the 'letter of the contract.' In light of the discussion, I propose:

> *Proposition 3: When contractual disagreements (over warranty claims) occur, an increase in interfirm managers automatically inferring opportunistic behavior on the part of their transacting partner, thereby enforcing the 'letter of the contract,' increases the likelihood of contract termination.*

The findings of the study show that transacting parties can have honest disagreements, and the primary reason for disagreements is errors in selection of decision factors due to reasons of bounded rationality in face of complexity of the underpinning complex technical system. In addition, contractual disagreements and contract termination can occur in the absence of opportunism, and the outcome is affected by whether managers act in the 'spirit of the contract' or apply the 'letter of the contract.'

2.6.4 Managerial Implications

Contracting for complex technical systems poses a significant governance challenge – contracts are unavoidably incomplete, often agents are dealing with ill-structured problems, erroneous selection of decision factors can occur leading to contractual disagreements. Under circumstances where problems arising during contracting for complex technical systems are not ‘sharply’ defined, the “sharp in by clear agreement; sharp out by performance” (Macneil, 1974: 738) approach, i.e., holding exchange partners to the letter of the contract seems unwise, and is likely to lead to contract termination. Yet, this approach of holding the transacting partner to the letter of contract seems to be common. The relevant question for managers is – is there a more prudent approach to contracting for complex technical systems such that the likelihood of contractual disagreements is low?

Consider Scenario 2, in which Indian Railways manager and OEM vendor managers, contracting for the complex technical system of the turbosupercharger, were unable to establish definitively their respective claims regarding the efficacy of engineering (assembly) specifications and performance outcomes for the component. Managers from both Indian Railways and the OEM vendor seemed to recognize the incompleteness of their knowledge. The goal – healthy locomotives – was in the joint interest of both Indian Railways and the OEM vendor. For Indian Railways, it meant higher availability of a key resource, which, in turn, meant lower warranty costs and a satisfied customer for the OEM vendor. Managers from both Indian Railways and the OEM vendor also recognized that they were dealing with an ill-structured problem. In jointly recognizing that they were confronted with an ill-structured problem, managers opened the possibility of negotiating a mutually acceptable solution, thereby avoiding an unnecessary protracted contractual disagreement and contract termination.

It is also in the interest of parties contracting for complex technical systems to increase their credibility by remedying actions taken based on their erroneous conclusions. For example, in

Scenario 5 pertaining to premature failure of liners, Indian Railways reinstated the downgraded OEM vendor to Tier 1 status. By reinstating the vendor to Tier 1 status, not only did Indian Railways signal to other prospective vendors in the industry that the railroad was not opportunistic, they also demonstrated the willingness to correct an honest mistake. Given that locomotive engines are known to catch fire, can cause immense damage, even loss of life, the decision made by Indian Railways to err (excessively) on the side of caution is quite understandable.

The above-described scenarios can serve as a template for parties involved in contracting for complex technical systems. To mitigate the possibility of an avoidable increase in friction, managers of transacting parties engaged in contracting for complex technical systems can mitigate protracted and unresolved contractual disagreements by viewing the contract as a flexible instrument for facilitating exchange, and not an instrument for inflicting undue damage. Viewing the contract as flexible instrument to enable economic exchange is congruent with relational contracting "when the problem [of contracting] is not one of honesty but one of reaching an agreement that both sides understand" (Macaulay, 1963: 58-59). The relational contracting approach does not suggest that opportunism is never at play, and safeguards be foregone, but only that opportunism should not be treated as a default premise to explain (unexpected) failures of the complex technical system being contacted for, but be evaluated on a case-by-case basis.

2.6.5 Limitations

In interpreting the findings of this study, some care needs to be exercised. First, the study maintains that caution ought to be applied in contract enforcement as managers might erroneously explain (unexpected) failures of complex technical by invoking opportunism. It does not suggest that opportunism is never in play, and not an important consideration for contract design. Second, this study pertains to the study of locomotives engines, and, and although it may be intuitive for managers to draw analogies to other industries such as the automobiles, such analogies may be flawed. For example, beyond the similar architecture, locomotive engines bear little resemblance to car engines in actual operation. A locomotive engine literally ‘redlines’ every single time it starts (when it is hauling freight), while most cars are not pushed to their extreme performance after every stop sign. Hence, the findings of the study ought not to be extrapolated without care. Third, the study pertains to the examination of contracts between a state-owned entity and private firms, and the findings may not map directly to cases of contracting where both contracting entities are private firms. Fourth, although there was no evidence of Indian Railways buyers dealing with the (OEM) vendors in a peremptory fashion (e.g., Williamson, 2008a), considerations of power affecting contracting cannot be definitively ruled out.

2.7. Conclusion

This study explores whether earnest buyers and sellers contracting for complex technical systems can have honest disagreements, and if so, what the antecedents are. In doing so, the study answered the call for research in OM examining the “day-to-day functioning” of operations and develop “implications” for TCE (Anand and Gray (2017: 5). The findings of the study indicate that in contracting for complex technical systems, bounded rationality affects the selection of decision factors, and when managers of transacting parties are confronted with ill-structured problems that have their origin in the technical failure of the underpinning complex technical system being contracted for, the likelihood of contractual disagreements increases. These disagreements can be resolved when the managers of transacting parties act in the ‘spirit of the contract.’ When managers of a transacting party assume that their transacting partner is opportunistic and enforce the ‘letter of the contract,’ the likelihood of contract termination increases. These findings suggest that, in practice, contractual disagreements and contract termination can occur as a result of honest disagreements.

Predicated on these findings, I suggest that managers are better served by treating opportunism as an empirical premise to be determined on a case-by-case basis. Mistaken attributions of opportunism by managers of transacting parties are an avoidable source of friction for contracting. The study also finds that agreement between managers of transacing parties predicated on flawed premises about the underpinning complex technical systems can be counter-productive. When managers encode these false premise about the complex technical system in their contract, they can arrive at unsatisfactory conclusions that are often inadequate to explain technical failures. While these premises (e.g., decomposability) can be convenient for managers to adopt, they affect how attention is directed when unexpected outcomes occur, and cause managers to overlook salient information.

Given the ubiquity of complex technical systems managers seek to manage, prudence requires managers to recognize that in dealing with technical failures, they are often confronted with ill-structured problems. These (unexpected) failures of the underpinning complex technical systems constitute an independent source of uncertainty for managers. When such technical failures (inevitably) occur, rather than defaulting to the explanation of opportunism by their transacting partner, managers of transacting parties can take the opportunity to jointly study these failures to expand their understanding of complex technical systems. In managing complex technical systems, managers are likely better served by trying to better understand why they disagree with their transacting partners when they do.

FIGURE 2.1

Adapted Warranty Record (sensitive material withheld)

S.no	1	2
Loco no	12269	12372
Rly	NWR	NFR
Shed	BGKT	SGUJ
Warranty Regn.No	(withheld)	(withheld)
Date of Regn.	02.01.12	03.01.12
DOC	01.01.10	04.02.11
DOF	31.10.11	22.11.11
Firm	(withheld)	(withheld)
Main Asm	Vehicle	Power Pack
Sub Asm	Cab Fan	TPU sensor
Item Sl,no	N/A	G2210
Qty	1	1
Complaint	Not working	Turbo
Status	Firm has sent to dlw for replacement	Closed
Closed Date	12-Mar	12-Mar
Cost	(withheld)	(withheld)
Part Number	18650016	17919
Cost recovered	(withheld)	-

Legend

Loco no: unique locomotive identification number

Rly: code for railway zone as designated by Indian Railways

Shed: code for maintenance depot

Warranty Regn.No: unique warranty registration number

Date of Regn.: date of registration of complaint

DOC: date the unit went into service

DOF: date of failure of unit specific to component

Firm: name of vendor/ supplier

Main Asm: main assembly

Sub Asm: sub assembly within main assembly

Item SI.no: unique item serial number of failed component

Qty: quantity of failed components

Complaint: brief summary of failure and location

Status: status of replacement claim

Closed date: date the claim was officially closed

Cost: cost of component in Indian Rupees

Part Number: unique part number of failed component

Cost Recovered: actual amount recovered from vendor

FIGURE 2.2

Data Mapping

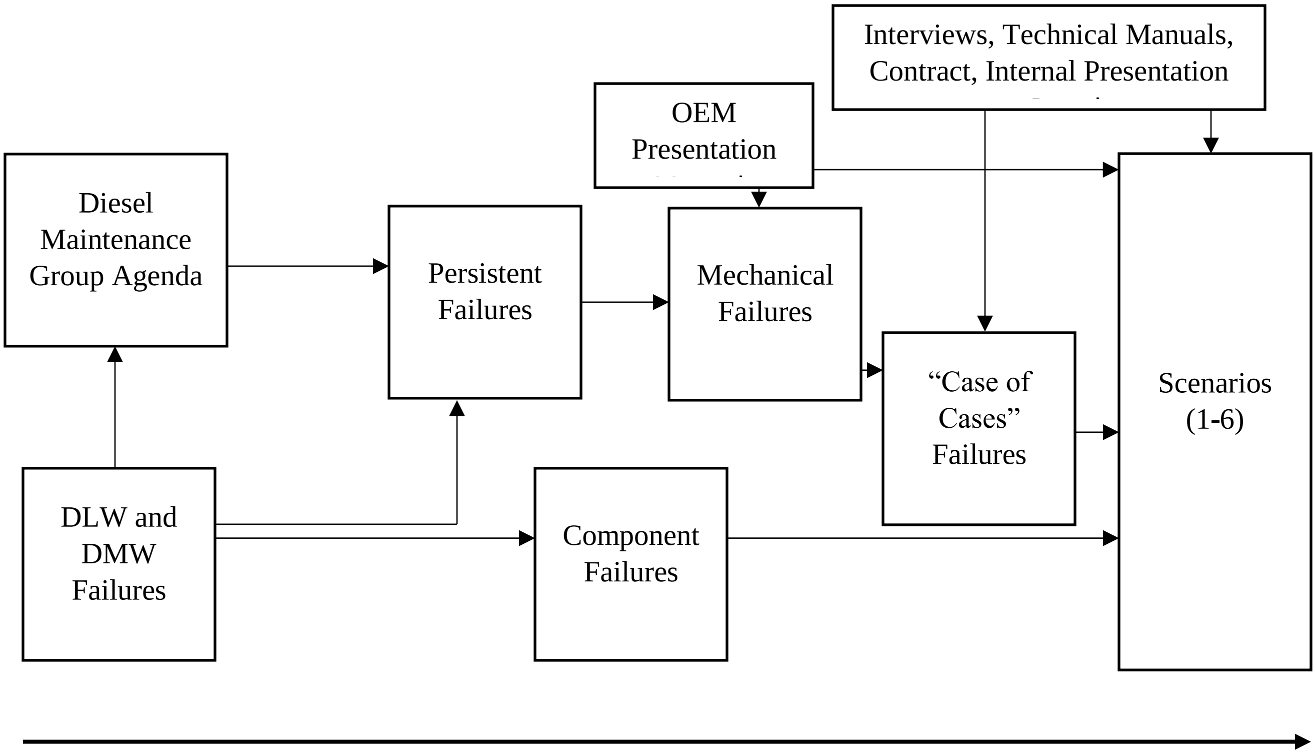

TABLE 2.1

Well-Structured and Ill-Structured Problems

Well-Structured Problems	Ill-Structured Problems
Scenario 1: Agreement between transacting parties	Scenario 2: Turbosupercharger assembly (disagreement) Scenario 3: Premature piston ring failure (disagreement) Scenario 4: Locomotive engine quality (disagreement) Scenario 5: Engine liner failures (termination)

Chapter 3

Contracting under uncertainty: The governance of high-reliability artifacts operating in the shadow of inevitable failure

Abstract

Concern with economic and technical failures has led scholars of economic organization and safety scientists to develop theories aimed at facilitating efficient economic exchange and safely managing complex systems. However, because of their differing foci, gaps in our understanding remain. In practice, managers must coordinate high-reliability artifacts, such as railroads, which contain governance structures, organizations, technical systems, and personnel interacting with each other. Failures of these so-called high-reliability artifacts are unpredictable, potentially catastrophic, and arguably inevitable. These failures often lead to contentious debates among contracting parties concerning root causes. To address this vexing contracting problem, I develop a framework that addresses managers' contractual choices to coordinate better these high-reliability artifacts, which operate in the shadow of inevitable failure. This study posits that mitigating failures of high-reliability artifacts are in the joint interest of the transacting parties. The developed framework examines factors that influence contracting effectiveness such as self-interest, opportunism, foresight, expertise, and complexity of contracts to help explain and predict the likelihood of artifact failures. I submit that the factors of differential hindsight, foresight and expertise are especially critical in the context of complex socio-technical systems. Aligning decision rights with expertise, contracts that empower those with greater local knowledge, and unanticipated outcomes in the technical system that receive subsequent managerial attention can lower the likelihood of systems failures. Finally, the domain of application of the developed framework, limitations, and possible future research directions are considered.

3.1. Introduction

Scholars recognize that failures of governance and of technical systems are ubiquitous, and constitute a vexing organizational problem. Such artifact failures not only can result in substantial economic loss, but also can (simultaneously) lead to substantial social costs, that can include fatalities (McMillan & Overall, 2017). In response to artifact failure, the extant literature has focused on mitigating contractual failures (Garicano & Rayo, 2016; Klein, *et al.* 2019), and making technologies safer (Dekker, Cilliers, & Hofmeyr, 2011; Reason, 1997; Weick & Sutcliffe, 2001).

Mitigating artifact failures via better contracting remains elusive, in part, due to silos within and across theoretical approaches (McMillan & Overall, 2017). Additionally, the *focus* of these different theories differ. For example, safety science focuses on managing complex technical systems (Dekker *et al.* 2011; Leveson, 2004). In contrast, transaction cost economics is primarily concerned with mitigating contractual hazards stemming from behavioral uncertainty (Crook, *et al*. 2013; Foss & Weber, 2016; Williamson, 1985). These different theories implicitly consider these different managerial choices to be independent. However, Farjoun notes that contributing factors to artifact failures can involve both "technical and organizational factor[s]" (2005: 27).

In the current study, I join theories from safety science and transaction cost economics to offer insights on the governance of high-reliability artifacts operating in the shadow of inevitable failure. Some parallels between the theories already exist. For example, to mitigate failures, the concept of "defenses" in the safety science literature (Reason, 1997) has its corresponding analogue of "safeguards" in the transaction cost economics literature (Williamson, 1985). The current study's unit of analysis is the artifact and focuses on artifact failures that are costly, unforeseeably, and arguably inevitable (Baum & Dahlin, 2007). The unit of analysis of *artifact* refers to synthetic systems designed to meet “human goals and purposes” (Simon, 1996: 3), an example being a railroad system designed to transport oil by rail in North America. This study considers *artifact in its entirety,* including not only considerations of economic exchange, but also the technical properties and operations of the complex socio-technical system. This study seeks to improve (organizational and technical) artifacts’ functions so that they are less likely to fail.

Under circumstances where avoiding failures are largely an unambiguous and undisputed goal shared by all stakeholders (Marschak & Radner, 1972), the key problem at hand becomes that of effective adaptation and coordination. To that end, I develop a framework concerning managers’ contractual choices to coordinate better high-reliability artifacts when there is unforeseeable uncertainty. Consequently, uncertainty in the behavior of the underpinning complex (technical) systems takes center stage. The developed framework examines factors that influence contracting effectiveness such as self-interest, opportunism, opportunism, foresight, expertise, and complexity of contracts to help explain and predict the likelihood of artifact failures.

3.2. Theoretical approaches to contracting in management

Even though the word *contract* often invokes legal connotations, in the case of business transactions, contracting is primarily a *managerial* or *organizational* concern reflecting a business consideration. Rubin states: “when a contract or transaction takes a certain form, the concern should not be with the legal form of the contract, but with the underlying business reality which makes the legal form desirable” (1993: xix). Accordingly, the management research literature on contracts focuses almost exclusively on the business and the ecosystem in which it is embedded, as well as the problems that the business managers seek to address. This study takes as a starting point the extant literature on transaction cost economics and its focus of attenuating bounded rationality problems and mitigating potential opportunism hazards (Williamson, 1985). I will then consider additional adaptation and coordination problems beyond incentive alignment.

Transaction cost economics considers contracting from an organizational and governance structure (rather than production function) perspective. Further, transaction cost economics views contracts as frameworks that provide order to an exchange relationship (Ketokivi, Mantere, & Cornelissen, 2017). Transaction cost economics considers all contracts as incomplete: it is impossible to incorporate all possible futures and events into “a comprehensive *ex ante* bargain” (Williamson, 1985: 28). The theory emphasizes the importance of foresight and contracting in its entirety by attempting *ex-ante* to devise safeguards (e.g., mediation and arbitration clauses, and vertical integration) to reduce *ex-post* transactional problems (Williamson, 1996). The main case hypothesis is to “align transactions (which differ in their attributes) with governance structures (which differ in their costs and competencies) in a discriminating (mainly transaction cost economizing) way” (Williamson, 1996: 46-47), and the critical attribute that determines the ensuing governance structure is asset specificity (Williamson, 1985, 1991). Transaction cost economics further posits that in the absence of proper safeguards, some contractual parties may choose to act opportunistically, and it is difficult to predict who will act opportunistically.

3.2.1. Some shortcomings of transaction cost economics theorizing in the context of high-reliability artifacts

Transaction cost economics focuses on the characteristics, objectives, and frailties of the transacting parties. The objectives of the transacting parties are, at best, partially congruent. While TCE acknowledges that uncertainty and complexity challenge the limits of our cognition (Williamson, 1975), it does not consider the differential impact of various types of uncertainty (e.g., demand, technical system, regulatory) on governance structure choice.

However, what if we shift focus from the transacting parties to the systems in which the contract pertains, which is the artifact (Simon, 1996)? Instead of taking the individual transaction as the unit of analysis, the current study considers interdependent and intertemporal transactions associated with the artifact in their entirety (Argyres & Liebeskind, 1999; Kang, Mahoney, & Tan, 2009). In particular, the current study considers governance of the complex *artifact in its entirety*, including the technical properties and operations of complex systems that occur within and across organizational boundaries. For example, consider the problem of safe oil transportation by rail in North America. In this case, the artifact of interest between the years 2008 and 2012 experienced a 28,000% increase in volume of highly inflammable crude oil transported (Therrien, *et al*. 2016: 265). In Simon's (1996: 5) terminology, this task constitutes an artifact that is influenced by technical and human systems; functions, goals, and adaptation; different legal jurisdictions; and by managerial decisions and actions. Because of the dynamic nature of the phenomenon of artifacts, there are limitations to invoking *ceteris paribus* conditions. Further, there are limitations to overlooking emergent outcomes arising from myriad interactions between human and technical systems, and the technical systems themselves.

I submit there are a number of key characteristics of artifacts that warrant attention, because these characteristics have profound governance (design) implications. The first characteristic is that artifacts tend to consist of interdependent subsystems that may be architecturally and organizationally independent but operationally, highly interdependent in that these subsystems interact in non-trivial ways. This point is crucial, because these non-trivial interactions give rise to emergent behavior, which are often associated with system failures (Perrow, 1984). Second, some *technical failures are arguably unavoidable*. Consider locomotive engine failures, which are unavoidable, unpredictable, and fail often enough to be a matter of grave concern. Because their assembly typically involves components supplied by multiple vendors, engine failures lead to non-trivial *ex-post* contracting (maladaptation and compensation allocation) problems (Bhardwaj & Ketokivi, 2016). Engine failures can lead to catastrophic disasters and massive economic costs – the fatal Lac-Megantic derailment being a recent tragic example in which 47 people lost their lives, and the operating railroad subsequently went bankrupt (Therrien *et al*. 2016). Further, a crucial but often overlooked factor is that often artifacts encompasses open systems (Thompson, 1967), where multiple actors, with varying expertise and belonging to different legal entities and jurisdictions, might interact unexpectedly leading to unanticipated outcomes. Current theorizing in contracting tends to underemphasize uncertainty that does not have its roots in human behavioral uncertainty (e.g., opportunism), and implicitly treats economic systems as if they were predictable and foreseeable enough to limit these type of uncertainties. Reductionism by parsimoniously focusing on different types of uncertainty leading to incomplete contracts and contractual hazards are workable in some managerial contexts. However, there are managerial problems in which a typology of types of uncertainty will matter. Consider bumper-to-bumper insurance coverage for car rentals – whether vehicular failure (damage) occurs due to incompetence, neglect, lack of care of the driver, or some other extraneous circumstance, the fee schedule can enfold the expected cost, and it does not matter very much *why* it occurred (assuming there was no collision involved). However, in the case of high-reliability artifacts, discerning the cause of failure is paramount. The assignment of massive liability damages and searching for solutions to mitigate the likelihood of future failures depend on locating the source of the component failures, which may occur due to poor quality of supply, poor design, unanticipated interaction, or some other technical reason.

3.2.2. Linking theories of economic organization and accidents

In artifacts beset with complexity, problems tend to be ill-structured (Simon, 1973) – determining the source of a failure due to a five-way interaction of components is an example. There is a tendency to try to establish root cause to facilitate the assignment of (financial) responsibility (MacDuffie, 1997). The human brain typically engages in *linear thinking*, i.e., assumes that failure is a process "that follows a chain of causal reasoning from a premise to a single outcome" (Dekker *et al*. 2011: 939). However, in practice, system failures are emergent outcomes that involved complex interactions and are not typically caused by a single factor (Perrow, 1984). Sometimes in these failures root causes cannot be identified (Dekker *et al*. 2011; Tucker & Edmondson, 2003). Moreover, these systems tend to exhibit safety drifts toward failure (Dekker, 2011; Hollnagel, 1993). Identifying the cause of failures depends on where the boundary is drawn. For example, while one might plausibly argue that the technical cause of the *Challenger* disaster can be traced to failed O-Rings, it is not reasonable to state that it was the *root cause*. Janis states: "[c]ommunication, as well as combustion, was responsible for the tragedy" (1991: 235). Further, the relationship between inputs and outputs cannot be precisely measured (Alchian & Demsetz, 1972), which hinders alignment of incentives, impedes development of adaptive mechanisms, and opens the door to opportunism (e.g., free-riding problem).

In addition, complex artifacts are often not well defined by the transacting parties (Farjoun & Starbuck, 2005). The operation of complex socio-technical systems also involves tradeoffs among various inputs and outputs (Reason, 1997). For example, safe transportation of oil by rail in North America is typically a mutually agreed-upon goal. That said, there might be substantial disagreements among transactional parties about acceptable private and social costs, as well as how to adapt better under uncertainty. The problem of allocating these costs is often exacerbated when managers holding particular decision rights do not possess the necessary information and the required expertise to make these decisions and adaptations (Farjoun & Starbuck, 2005). Worse, when failures occur, these managers may be unable to quantify fully the overall damage these failures cause (e.g. Deepwater Horizon). The institutional environment in which these managers operate may not be amenable to change and shapes response to these failures (Vaughan, 2005). Boundedly rational agents not only do not possess the foresight required to be able to evaluate fully the implications of their decisions, especially when the environment is complex (Augier & March, 2008), but also they may not even be able to determine the adaptive effectiveness of their previous decisions due to the complexity of the interdependent problems they encounter.

In large-scale artifacts such as those involving transportation of crude oil, the feasibility of modifying organizational boundaries to find the comparatively superior governance choice may be severely limited. The crux of the problem is that even recurring, complex transactions beset with asset specificity and uncertainty cannot always be internalized, even though that is what TCE would typically prescribe. For example, vertical integration is not necessarily feasible especially if oil is transported by a state-owned railroad (e.g., Indian Railways). Considerations of national security alone rule out the purchase.

In summary, what is missing is an approach for contracting in situations that are characterized by ill-defined problems, which arise due to emergent system behaviors and unclear goals (e.g., safety vis-à-vis cost considerations). Moreover, conventional measures such as vertical integration as a safeguarding mechanism of better aligning incentives and internalizing the risk to a single organization are constrained. In addition, some form of system failure is inevitable and has the potential to inflict massive damage. In such contexts, which often involve multiple transacting parties, in the case of a failure, even when the cause of failure is due to decisions by an outsourced contractual party, both parties may still legally be jointly accountable for damages, as was the case with Lac Megantic.[1] These contractual problems can be found in many different empirical settings, such as energy production, public transportation, space exploration, and many industrial contexts where large OEMs supply large-scale systems for their clients, who in turn operate complex production or service systems. I maintain that for these high-reliability artifacts, in general, and for railroads specifically, that it warranted to posit that avoiding failures is in the joint interest of all stakeholders, and consequently, we can claim a common objective.

To provide remediable solution that improve safety in high-reliability artifacts, such as railroads, I examine the consequences of cognitively limited managers coping with complex behavioral and environmental factors. These factors are comprised of self-interest, opportunism, bounded rationality, limited foresight, differential expertise, what information is used and to what ends, the implications of not being able to identify a root cause, and decision-making under uncertainty. It should be noted that some failures that emerge from system behaviors are unknown and some perhaps are unknowable. Recognition of such deep uncertainty can be connected to Williamson's (1985, 1996) concept of remediableness and challenges in improving (governance) systems. Williamson elaborates: "Some opportunities for improvements will not be perceived at all; some mistakes will be recognized only after that fact. Any 'shortfalls' due to misperception or mistake *will not be remediable* by supplanting contract by vertical integration" (1985: 66, footnote

---

[1] See also, Case No. 3:15-cv-13328 Sigman et al. v. CSX Corporation et al.

1, emphasis in original). Remediableness is a key criterion of the current study. For example, conventional governance solutions such as vertical integration may not be precluded, and in such cases, only feasible alternatives will qualify for comparative assessments.

3.3. Contracting under uncertainty

In this section, I introduce a framework contracting under uncertainty for the management of high- reliability artifacts that inevitably experience some failures that cannot be predicted, and might inflict massive damage. These artifacts are synthetic systems that are designed to "meet human goals and purposes" (Simon, 1996: 3), and incorporate cognitively limited agents in complex relationships with the system, environment, each other, and embedded in a specific institutional setting. Moreover, these artifacts cast 'weak empirical shadows' in that their inner workings and interactions are not directly manifested in their behavior and remain partially hidden, and reside beyond managers' comprehensive understanding and control, and reflect their complexity. Decision-making under such uncertainty can entail substantial ambiguity of identifying and selecting relevant premises and their corresponding appropriate inferences associated with observed outcomes. Consequently, the relationships between inputs and outputs are not clearly discernable or (precisely) quantifiable. Managers formulate judgments about the future course of events, which may not have a firm empirical basis (Knight, 1921). Additionally, there are multiple diverging, even conflicting goals among contractual parties, some of which cannot be precisely operationalized, and hence, can be characterized as "ill-structured problems" (Simon, 1973).

Managerial knowledge might consist of how artifacts (systems) work, but rarely on the full set of ways that artifacts could fail. Moreover, managerial knowledge is distributed throughout the organization and individual managers have local knowledge of the system related to their own expertise and decision-making responsibility. Thus, the very complexity of artifacts often exceeds "[managers] ability to manage them" (Farjoun & Starbuck, 2005: 5). Design features have the propensity to cause unforeseen and invisible reverberations throughout the artifacts, which may only become evident during their operation (Dekker, *et al.* 2011). Not only do managers have bounded rationality in which the complexity/uncertainty of their contractual problems exceed their capacity to consider all possibilities ex ante (Simon, 1957), but contrary to Williamson's (1996) strong-form premise concerning managerial foresight, managers are also posited here to have *boundedly foresighted.* Thus, managers may not fully take into account the incompleteness of their contract (Mayer & Argyres, 2004), and thus exclude necessary safeguards (e.g., mediation and arbitration clauses for when unexpected conflicts occur). Here, the ability to foresee *behavioral* uncertainty and enfold appropriate safeguards at the contracting stage has limited utility.

Further, because bounded rationality puts limits on managers' knowledge, abilities to acquire expertise, and processing capabilities (Simon, 1957: 198), they may be unwittingly mistaken in their classification of the cause of artifact failure, which impedes the detection and correction of error. The central problem for managing such artifacts is adapting to uncertainty stemming from (1) incompletely understood emergent behavior of the underpinning complex technical system; and (2) unknown, unpredictable, and unforeseeable future states of the world resulting from (unexpected) interactions within and across the interface that can cause failures. To tackle this problem, I develop a framework in which I enfold considerations of uncertainty that have their origin not only in human behavior, but also in the complexity intrinsic to the artifact (Hayek, 1945, 1979; Knight, 1921).

3.3.1. Developing the framework

In the current section, I develop a framework for the governance of high-reliability artifacts that operate under the shadow of inevitable failure (e.g., railroads, space exploration, and power plants). The critical underpinning assumptions are that the safety goal of mitigating failures is in the joint interest of transacting parties, and that (some) failures are unpredictable, unpreventable, may cause massive damage, and may occur with a frequency which is high enough to constitute a managerial problem. Under these circumstances, I consider implications of behavioral premises, such as opportunism, limited cognition, expertise, and contract complexity.

3.3.1.1 Opportunism

Economic theories of contracting tend to focus on attenuating uncertainty stemming from opportunistic behavior of contracting parties. Forms of opportunism include *ex-ante* asymmetric information problems (e.g., adverse selection), *ex-post* asymmetric information problems (e.g., moral hazard), as well as attempts to renegotiate contracts, the so-called economic holdup problem (Arrow, 1985; Williamson, 1975). The usual approach is that since some contractual parties may behave opportunistically, it is prudent to design contracts that anticipate potential contractual hazards and provide necessary safeguards to mitigate these problems, such as vertical integration and mutual sunk cost commitments (Williamson, 1985).

However, such *game-theoretic* uncertainty captured by the opportunism construct is not the only type of uncertainty in which managers must respond. Another important type of uncertainty is decision-theoretic rather than game-theoretic. Namely, decision-makes cannot fully comprehend and anticipate emergent outcomes of underpinning technical systems and interactions at the socio-technical interface in complex artifacts (Clearfield & Tilcsik, 2018; Perrow, 1984).

Looking back from an outcome does not always enable managers to trace a path to a singular origin of an artifact failure. The sheer number and diversity of interactions within complex artifacts preclude managers from definitively identifying the root cause of the artifact failure (Dekker *et al.* 2011; Waldrop, 1993). Further, any analysis requires a decomposition of an artifact into more manageable subsystems (Hayek, 1964; Simon, 1996), but there is the possibility that the decomposition initiated may make it virtually impossible to isolate the root cause if the managers place key contributing (interdependent) factors of the artifact failure in different subsystems. The limitations of managerial ability to examine past events to comprehend fully the mechanisms underpinning artifact failure, I refer to as "*bounded hindsight*." Put differently, *bounded hindsight* is a cognitive constraint that precludes managers from effectively partitioning artifacts for analysis and establishing (unambiguous) causal chains of events leading up to an artifact failure.

However, since managers must draw a boundary to conduct a meaningful analysis, they face the possibility of inadvertently discarding essential information and thereby defeating the very purpose of the analysis. It follows that managers might begin their analysis on false grounds (e.g., decomposing the problem incorrectly) and being led to false conclusions. Furthermore, root cause analysis is often used to assign financial responsibility, and may not result in mitigating future failures, which is prescriptively the objective of root cause analysis (MacDuffie, 1997). For example, the extant literature provides little evidence that the prescriptive objective of root cause analysis is implemented effectively in the field of medicine (Wu, Lipshutz, & Pronovost, 2008).

Further, the extant literature on contracting, because it does not take into account that managers have bounded hindsight, may not fully capture the nuances that managers face in the context of artifact failure. Consider a scenario where *after* the contractual parties engage in an economic exchange, an unanticipated disturbance (event) occurs. A contractual party with foresight will seek to *ex ante* incorporate safeguards against the perceived cause of the artifact failure. If managers attribute the primary source of the artifact failure is opportunistic behavior the governance mechanisms will be chosen to mitigate the perceived problem. Transacting parties would then be expected to develop the capability to design better contracts (e.g., Mayer & Argyres, 2004). Nevertheless, regardless of the course of action chosen, identifying an incorrect root cause problem diverts attention away from learning about actual system behavior. Even when attention is directed towards learning about the relevant problem, it may not occur due to the inherent complexity of the artifact (e.g., Starbuck & Farjoun, 2005), the scarcity of failures of this type (March, Sproull, & Tamuz, 1991), or both. Any disturbance that occurs could be the result of an accidentally incorporated design feature or an unexpected inter-action that remains hidden. Alternately, unanticipated interactions, which could not possibly have been foreseen by cognitively limited managers, may be foreshadowing the possibility of an impending failure rather than being the result of opportunistic behavior.

When managers default to attributing opportunism to unanticipated events that can be explained by incomplete understanding of the artifact, they reinforce a misplaced confidence in their ability to understand and manage complex artifacts. It is when managers are confident of the possible outcomes they are most surprised when their expectations are proven incorrect. Under these conditions, the inability of a transacting partner to deliver on its promise can be misconstrued as opportunism. In the context of high-reliability organizations, unwarranted *mistrust* is particularly problematic. Thus, opportunism is not *necessarily* present, and can be incorrectly invoked. Prudence requires a thorough examination of artifact failure, which can be stymied by unwarranted assumptions of opportunistic behavior. While managers' limited cognitive abilities impede their ability to explain every unanticipated outcome that may occur in any economic transaction, presuming opportunism can impede their understanding of the system.

Critically, because opportunism is a *sufficient* condition to explain unanticipated outcomes, it can act as a *de facto* factual premise (Simon, 1976) for an organization. Such reasoning can severely affect and distort managers' understanding of the artifact, undermining trust and thus increasing the overall cost of contracting (Arrow, 1969). It can also lead managers to develop improper contractual safeguards and technical system defenses, which might fail to prevent or reduce future artifact failures. Managers' cognitive limits conjoined with opportunism is a vexing problem for contract design and execution. The concern highlighted here is that managers who rely on opportunism as the default explanation for each and every artifact failure will impede learning from the many cases in which the contributing factors for artifact failure are due to cognitive limits such as managers' bounded rationality and/or bounded foresight, and which are not identified as such. This mis-attribution error impedes managers from investigating other causes that may satisfactorily explain artifact failure, and the managers' (mis-)attributions of opportunism may even cause the contractual relationships to fracture, thereby foregoing mutual gains. To counteract this contracting problem, I propose:

> *Proposition 1: Transaction parties that increase their use of opportunism as the default explanation for artifact failure decrease their likelihood of learning relative to those searching for explanations rooted in non-behavioral (technological) uncertainty.*

3.3.1.2 Self-Interest

Self-interest is the foundation on which voluntary exchange rests (Smith, 1776). In modern economic exchange, the institutional environment, as for example in publicly traded firms

(Ketokivi & Mahoney, 2016), typically prescribes self-interest of transactional parties on behalf of their organizations. In typical situations of voluntary economic exchange, the objectives of transacting organizations partially overlap, which facilitates exchange relationships. Exchange partners anticipate that their transactions are a positive sum-game. However, circumstances may arise where, *ex post* an exchange partner may be worse off (e.g., if there is economic holdup). Current theories of contracts recognize this possibility and prescribe means to mitigate holdup problem (e.g., designing safeguards). The usual theoretical premise here is that the ex-post position of an exchange partner can be worsened because of opportunistic behavior, and contractual parties should take necessary measures such as vertical integration to protect their (mutual) self-interest (e.g., Klein, Crawford, & Alchian, 1978). However, this approach ignores the possibility of failure stemming from unforeseen emergent properties of complex socio-technical systems (which can leave one or both contractual parties worse off).

Before contractual parties agree to an exchange, it may not be possible to assign fully responsibility and costs for future states of the world (Coase, 1937; Williamson, 1993a). Forward-looking, cognitively limited managers cannot foresee every outcome that can occur. Additionally, for complex artifacts, the relationship between inputs and outputs is not always fully separable or (precisely) quantifiable. In the absence of clear causal chains, non-additive interactions, and difficulties with measurement, act as barriers to assigning responsibility and costs. Effective contracting be impeded by self-interest (with guile) and conditions of information asymmetry (Williamson, Watcher, & Harris, 1975). Even if it were possible to do so, in cases where the impact of failure spills over beyond the transacting parties, courts are not sympathetic to legal arguments that attempt to absolve individual parties of responsibility.

Under these conditions, I submit that mitigating failure is a *joint problem* that requires joint effort. This outcome occurs because unless it is treated as a joint problem requiring cooperation, any adaptation that ensues is likely to be ineffective. More troublingly, self-interest may prompt contractual parties to focus on developing more efficient *ex post* adaptive mechanisms in their contracts (e.g., Argyres & Mayer, 2007) at the expense of trying to prevent failures. Since managerial attention is a limited resource, a narrow focus on contracting severely limits learning about the artifact itself. While such an approach is justified in many circumstances, in the context of artifacts that have the potential to cause massive damage, it may be detrimental. Contractual parties acting in their self-interest, attempting to limit the fallout of economic failure may neglect to, or choose not to, share pertinent and vital information that mitigate artifact failure.

I suggest transacting parties are better served by making (mutual) credible commitments to mitigate failures. Such an approach is in the joint interest of transacting parties where "common, not divergent, interests and beliefs" prevail (Marschak & Radner, 1974: 4). Rather than employing the prism of narrow self-interest, transacting parties can seek to agree on how they can jointly pursue their goal of 'improving' the artifact. An advantage of this approach is that it goes beyond the narrow interests of particular organizations. For the approach to be successful, the transacting parties must perceive each other as being credible, absent which nefarious intent might be imputed. It is therefore incumbent on the contractual parties to build their credibility (see Williamson, 1996: 134-142 for discussion on the subject). Prevention of artifact failure is a joint undertaking.

While the overall interests of transacting parties may only be partially congruent, within the context of preventing breakdowns and catastrophic disasters, these goals do fully align. This outcome is particularly applicable for high-reliability artifacts where the cost of failures may be devastating and exceed gains that might accrue due to narrow formulations and differential goal pursuit by individual transacting parties. Organizations in these industries (e.g., nuclear energy, and public transport) tend to be heavily regulated, and failures can lead to criminal proceedings.

> *Proposition 2: Transacting parties with governance structures designed to promote their joint interest in achieving safety will be more likely to mitigate artifact failures relative to transacting parties without governance structures designed to promote their joint interests.*

3.3.1.3 Bounded Foresight

Williamson accords managers considerable foresight in recognizing their own bounded rationality and posits "incomplete contracting in its entirety" (1996: 9). It is critical to note that the focus of economic theories of contracting is mitigating behavioral uncertainty (opportunism) by contracting parties, as well as positing managerial foresight incorporate safeguards (e.g., mediation and arbitration mechanisms) to attenuate contractual hazards. The current study not only considers these relevant dimensions of contracting, but also considers uncertainty originating from the under-pinning complex socio-technical system, which remains unexamined.

To deal with complexity, Simon maintains that artifacts be decomposed such that most subsystems "are only weakly connected with other [subsystems]," and serves an analytical purpose as well (1996: 209). There can be an enormously large number of ways to partition a complex system and thus determining how artifacts can be effectively partitioned to be "nearly decomposable" is a non-trivial challenge (Bhardwaj & Ketokivi, 2016; Hayek, 1964). Because the system can reveal unexpected interactions and thus be less decomposable than originally imagined is why we observe unanticipated emergent behavior. Theories of accidents as well as complexity theory explain why emergent outcomes are unpredictable and can be qualitatively different from expectations (Waldrop, 1993; Dekker *et a*l. 2011; Weick & Sutcliffe, 2001).

In transaction cost economics, (unanticipated) events or disturbances are classified based on being inconsequential, consequential, and highly consequential, depending on their deviations from efficiency and accompanying adjustment costs. Adaptive responses are made by comparing expected costs of adaptation with expected gains from adaptation (Williamson, 1991). Because of the very high adjustment costs resulting from highly consequential disturbances, it is critical that managers design safeguards that reduce the likelihood of recurrence(s).

Much of the literature concerning contract design neglects two important consequences. First, the opportunity to learn about the potential non-behavioral (e.g., technical) causes of the disturbance are neglected in favor of exclusively focusing on developing adaptive mechanisms that attenuate behavioral uncertainty (e.g., opportunism). Second, because managers' attention is a limited resource (Ocasio, 1997), the possibility of ignoring a critical disturbance that is currently not of large economic magnitude increases. Yet in complex artifacts beset with non-linear feedback loops, so-classified 'minor' events can lead to massive failures. For example, prior to the horrific crash of the Challenger shuttle, NASA was aware of a host of problems, but chose to label many as non-critical maintenance issues (Feynman, 1986). While this approach economizes on organizational resources and managers' cognition in coordinating and contracting for high-reliability artifacts, it can highly jeopardize the jointly held mission of safety by contractual parties.

Because of these concerns, I note a theoretical gap in the extant literature that needs to be considered. Specifically, I suggest amending Williamson's (1996) *incomplete contracting in its entirety* to recognize that managers typically do not exhibit perfect foresight (see, e.g., Mayer and Argyres, 2004). Thus, I introduce the constructs of "*bounded foresight*" to capture the fact that managers are often unable to imagine how artifacts might fail, and consequently *do not* provide all the necessary feasible safeguards to mitigate artifact failure.

Imagining how artifacts might fail should not be taken as a purely exogenous variable. Managerial training can channel their attention to considering exceptions. Relatedly, Marschak and Radner suggested that decision makers focus on variables that take on “exceptional values” in the tails of the distribution for those variables (1972: 206-207). I concur, and add that “disciplined imagination” (Weick, 1989) by managers can also suggest “omitted variables” for consideration that are relevant for mitigating artifact failure, which facilitates emergent adaptation. Because the causal links between individual variables and outcomes are often opaque, emergent behavior is the appropriate focus of managerial attention.

The current study suggests that mangers treat these “exceptional values” and “omitted variables” as indicators of systems behavior that provide insights into why artifact failure might occur, and not to simply classify these indicators as anomalies (“bugs”) of the artifact, which will then run the risk of being ignored. Greater managerial habits of “systems thinking” will lead managers to be less inclined to ignore anomalies. Indeed, these anomalies will be taken seriously as potentially foreshadowing future artifact failure, which provide opportunities to learn and adapt. Note that focusing on anomalies need not be that costly to implement. For example, prior to the fatal Lac-Megantic derailment, the locomotive engine exhibited signs of distress, which were discounted as being significant (Campbell, 2018). Directing attention to the anomaly of the smoking locomotive spewing oil droplets from its exhaust subsystem would have caused a few hours delay of one train. To mitigate artifact failures and facilitate *ex-post* adaptation, contracts can be designed or modified such that anomalies associated with particular (safety-critical) subsystems trigger joint investigations into their cause. Therefore, I propose:

> *Proposition 3: Transacting parties that increasingly focus their attention on anomalies regardless of their economic magnitude are more likely to mitigate artifact failures relative to transacting parties that default to allocating attention only to failures of high economic magnitude.*

3.3.1.4 Bounded Rationality

Current theories of contracting tend to treat bounded rationality in a background sort of way, and usually invoke it to explain unaccounted for contingencies and incompleteness of contracts (Foss, 2003; Foss & Weber, 2016). For transaction cost economics, how the managers' limits of bounded rationality are reached is not relevant, computational capacity might be exhausted by a large decision tree (Williamson, 1975: 22-23). In contrast, I suggest that it is critical to understand how managers' cognitive bounds are reached. For example, bounded rationality can manifest itself in errors of classification. Because managers are boundedly rational, and organizations supply them with premises to economize on bounded rationality, there is a tendency of managers to fit anomalies into pre-existing 'buckets' or schemas that managers possess (Farjoun, 2005). In such cases, anomalies can either be glossed over as human agents consider them unique, or if they occur repeatedly and are predictable, become normalized (Vaughan, 2005). This tendency can be regarded as a confirmation bias, whereby evidence is interpreted in a manner that is "partial to existing beliefs, expectations, or hypothesis in hand" (Nickerson, 1998: 175; Wason, 1960). The contracting problem is precisely how to 'manage' high-reliability artifacts by encouraging managerial learning in systems that they do not fully comprehend, to mitigate artifact failure.

The contracting problem across a system, which can encompass many firms, can be addressed by generating near and hypothetical histories (March *et al.* 1991) concerning anomalies, as well as by involving organizational personnel with differing expertise. Doing so is (more) likely to generate richer multi-faceted narratives to plausibly explain these anomalies. Crucially, managers have a tendency to avoid even the recognition of uncertainty and exhibit a quest for certainty. Such a tendency, while perhaps psychologically comforting for many managers, is likely to lead managers to normalize anomalies by force fitting them into existing classification schemes (Weick, 2005). This approach can be attenuated with the use of simulations and experiments (Sterman, 1994). Alternative histories and narratives can be interchangeably substituted and evaluated in their ability to plausibly explain such anomalies. Rather than accepting explanations for artifact failure within the managers' current schema, which run the risk of being too simplistic and lacking in precision, transacting parties that jointly seek multiple explanations that challenge their current schemas are more likely to garner a better understanding of the underpinning mechanisms of artifact failure. While the process may consume (cognitive) resources, preferring expediency and prematurely accepting explanations may be counterproductive for mitigating artifact failures that in many cases are associated with human casualties. This reasoning leads to the following proposition:

> *Proposition 4: Transacting parties that jointly seek (multiple) alternate explanations for anomalies beyond their current schema increases their likelihood of mitigating artifact failures relative to contractual parties that default to categorizing artifact failures into their pre-existing organizational schemas.*

3.3.1.5 Expertise

In contracting for artifacts, the division of managers' knowledge poses a significant challenge for effective adaption (Hayek, 1937; Hoopes & Postrel, 1999). One governance approach for attenuating the problem of interdependence (of knowledge and expertise) under conditions of uncertainty and complexity is vertical integration, but this managerial choice is not always available. Even when feasible, vertical integration is typically a centralized approach (Milgrom & Roberts, 1987), and has the disadvantage of impeding autonomy and emergent ways of addressing artifact failure due to bureaucratic distortions (Hayek, 1945). An alternate approach seeking to capture benefits of greater autonomy is to contract for expertise (Mayer & Argyres, 2004). However, there is a tradeoff. Specifically, the benefits of specialization of these autonomous firms have their limits because interdependencies that might be considered in the vertically integrated firms may become blinds spots at the interface of autonomous organizations, where it is either unclear what the interdependencies are, or there are contractual difficulties on how to handle such interdependencies. Aligning expertise with decision rights is a challenge, and often situations arise in which managers assigned decision-making authority are not the managers most qualified to make those decisions (Farjoun & Starbuck, 2005: 13-14).

To overcome the contracting problem of effectively partitioning the decision rights to match appropriately managerial authority with expertise (and local knowledge). There is a need then for managers to focus on systems thinking to reduce artifact failure. Managerial attention must be given to *joint* problem solving and learning, and in partitioning and assigning roles and responsibilities based on expertise. Importantly, the mutual goal of mitigating artifact failures must be clearly communicated, accepted, and adhered to by contracting parties, even if this mutual goal leads to short-term losses for one or more of the contracting parties (Kreps, 1996). The mutual goal of mitigating failures acts as a focal point (Schelling, 1960).

In comparing organizational systems for mitigating artifact failure, this study suggests that proposed detailed contracts and design rules need to be evaluated comparatively. Specifically, comparatively efficient organizational governance structures to mitigate artifact failure are those that more effectively identify experts who are assigned the decision rights (Argyres & Mayer, 2007). Because experts from *across* organizations are involved in the design process, ambiguities of language and instances that require conflict resolution are likely to be reduced over time. Integrating various aspects of operation and governance, and enfolding them in the design requires coordination of interdisciplinary experts (Hoopes & Postrel, 1999). When such integration requires shared language among an interdisciplinary group of specialized experts, repeated interactions are key to achieve the efficiencies inherent in generating and diffusing a common code (Arrow, 1974), thereby attenuating problems typically associated with misunderstanding and limits of language. This approach will likely elevate immediate costs, which will nonetheless often be justified in terms of the benefits of achieving the mutual goal of attenuating artifact failure. Hence, I propose:

> *Proposition 5: Organizational systems that better match (inter-disciplinary) expertise with responsibility and decision rights for contract design increase their likelihood of mitigating artifact failure.*

3.3.1.6 Contract complexity

For artifacts, adaptation to uncertainty is the central economic problem, and can be resolved either by centralization or decentralization (Hayek, 1945). At one end is a single authority making every decision, on the other resides complete atomization. Contracts reside along the continuum, and can act as an adjustable framework to facilitate adaptation (Llewellyn, 1931). These contracts vary from coarse-grained to fine-grained, with the latter being more complex, and entailing greater specifications of details, such as processes and obligations (Poppo & Zenger, 2002). Typically, economic and legal theories of contracting prescribe complex contracts for complex transactions (Argyres & Mayer, 2007; Poppo & Zenger, 2002). Models of contracting also predict increased contract complexity in the face of uncertainty (Eggleston, Posner, & Zeckhauser, 2000).

The extent and specificity of rules have a direct bearing on adaptation. The more fine-grained the contracts, the more specific are their terms, the more detailed they are, the more it approaches coordination by centralized authority with the autonomy of the local managers being limited. Coarse-grained contracts accord more discretion to managers, who, based on their local knowledge, have greater autonomy to respond to events as events unfold (Hayek, 1945), albeit remaining within the bounds of prior contractual agreements. Coarse-grained contracts may also be favored over fine-grained contracts because the latter increase the likelihood of making errors as local managers may misapply a rule and make errors in classification especially where outcomes that obtain are not perfectly measurable or are ambiguous in nature.

Further, rather than economizing on bounded rationality, complex contracts add to the cognitive load by increasing the amount of overall computation required. There are several reasons for this cognitive strain. First, because managers are boundedly foresighted, their ability to foresee contingencies for which adaptive responses need to be codified is questionable at best. For any outcome that occurs, a corresponding rule might not have been formulated causing the system to ‘freeze’ as managers search for proximal rules. Managers may also find themselves in situations where adaptive options are constrained by predefined commitments. Second, as the degree of granularity increases, complexity increases, taxing the cognitive and processing capacities of local managers. Relatedly, as the number of hierarchical levels of the contract increases, the number of acts of classification required by boundedly rational managers to reach the applicable abstract rule increases. Not only does this classification exercise require more time, but it also increases the branches of the decision tree and the number of nodes at which an error in judgment may occur. Third, increasing contracting complexity increases the number of possible interactions between different rules, elevating the likelihood of undesired consequences and mutually incompatible rules. Fourth, designing detailed contracts is both costly and time-consuming. Correspondingly, monitoring becomes difficult and costly. Additionally, because (cognitive) resources are limited, developing fine-grained contractual responses for anomalies that might not occur is akin to manages crossing all the bridges that they might encounter, which can be highly inefficient, and comes at the expense of attempting to understand the functioning of the artifact and mitigating artifact failures. Finally, some anomalies can occur, which may not be possible to categorize, *ex ante*. Put differently they involve Knightian uncertainty (Knight, 1921). Hence, complex contracts are likely to impede adaptation. Therefore, I propose:

> *Proposition 6: Organizational systems that use complex contracts are less adaptive in complex systems, which increases the likelihood of artifact failure.*

There is, of course, a limit to which a contract can be simplified. This reasoning leads to the conclusion that contracts that are too simple, and contracts that are too complex are fraught with difficulties that can contribute to artifact failure. What is the Goldilocks principle? This study suggest that as the overall complexity of the artifact increases, a point will be reached in which there are too many potential bridges that need crossing that must be considered. The advantages of hierarchical firms take on greater significance. In essence, the firm enables managers to use the authority relationship and to coordinate by fiat – within a “zone of acceptance” (Simon, 1976) -- to cross the bridges as they come to them in a more adaptive, sequential manner relative to other contract alternatives (Williamson, 1985: 89). Note, while Williamson concurs with this idea, the current study is distinct because it suggests that such considerations provide a *technical systems rationale for vertical integration*, which focuses on addressing non-behavioral uncertainties.

3.3.2. Testing the propositions

The propositions developed in the current study are applicable for managing complex artifacts (e.g., railroads, space exploration), and require an examination of artifacts comprising multiple organizations, actors, and technical systems engaged in a web of interactions (Campbell, 2018; Starbuck & Farjoun, 2005) that experience failure. The dataset should include contracts, records of failures, ensuing actions taken by contracting parties, their observations and comments on failures, correspondence between transacting parties, and notes and memos of meetings held between transacting parties. There are at least two broad categories of artifact failures that qualify for relevant areas of empirical investigation to test the developed theory in the current study, the first being (massive) technological failures that lead to public investigations by outside (federal) agencies (Campbell, 2018; Farjoun & Starbuck, 2005). The second category contains contractual failures of public utilities and stated-owned public transportation systems (Bhardwaj & Ketokivi, 2016) where information can be freely obtained on request.

Consider Proposition 1 concerning opportunism. The contracting parties will likely elevate expedited resolution of unanticipated outcomes over understanding because, absent massive failure, incentives favor continuing operations with current routines. The incentives are embedded within the warranty contracts such that failures automatically trigger a warranty request. Further, bounded hindsight impedes managers from identifying underpinning reasons for artifact failures. The default premise typically held by buyers, especially when vendors are vetted extensively for quality due to safety considerations, is that the transacting partner supplied lower-quality components. Similarly, consider Proposition 2 concerning self-interest, we can expect to see transactions in which large buyers deal with smaller suppliers in a peremptory fashion (Williamson, 2008a). Transacting parties and organizations that are not implicitly cognizant of bounded foresight, considered in Proposition 3, will likely not safeguard against inexpensive and minor failures (Feynman, 1986), as well as those that are safety critical. Managers' cognitive limitations, as discussed in Proposition 4, will result in mis-classifications, reluctance, or outright refusal to seek and to consider explanations of (component) failures that do not readily fit existing organizational schemas for anomalies (Weick, 2005). We are also likely to find misalignment of decision rights and expertise (Farjoun, 2005), which Proposition 5 suggests. Finally, with respect to Proposition 6, contract complexity, we can expect to find the presence of detailed bureaucratic procedures across silos within organizations as well as across organizational boundaries (Clearfield & Tilcsik, 2018). An additional implication of Proposition 6 is that we can expect to observe firms vertically integrating in the face of complexity despite no evidence of opportunism.

3.3.3 Application

The contracting approach developed in the current study emphasizes expertise, and treats coordination and cooperation to achieve the mutual goal of safety, which thus elevates mitigating failure over cost efficiency in the contexts of high-reliability artifacts. Indeed, *the unit of analysis* in this study can be considered as *the problem* of managing in the context of high-reliability artifacts (Nickerson & Zenger, 2004). In these contexts (e.g., railroads, space exploration), the institutional environment typically prescribes a certain level of information transparency, and the associated fall-out of failure can be very costly, or even fatal (Therrien *et al*. 2016). Further, there are legal precedents that establish joint responsibility of failure. The accompanying liabilities of catastrophic failures may lead to federal investigations, and even criminal proceedings against an organization (Dekker, 2011). The possibility of these outcomes tends to align the safety mission of the transacting parties.

The developed framework for contracting under uncertainty becomes increasingly appropriate as high-reliability artifacts can inflict massive damage when there are artifact failures. This study enfolds considerations of non-behavioral complexity, and recognizes that managers are boundedly hindsighted and boundedly foresighted, while retaining the traditional assumptions found in transaction cost economics of bounded rationality, the possibility of opportunism, and prescribing the design of safeguards (Williamson, 1996).

3.3.4 Opportunities for future research

The current study offers several avenues for future research. First, my developed framework can examine high-reliability artifacts failures for evidence of governance structures that are unsuccessful in rewarding the joint goal of mitigating these failures and instead focuses on penalizing any failures that might occur (Chikudate, 2009; Conklin, 2016). Such cases can shed new light on why artifact failure occurs despite the mutual goal of attenuating such failure. Thus, I suggest future studies that consider unexpected interactions and unintended consequences of governance structure choice. Second, new cases studies would enable us to triangulate to offer richer accounts of high-reliability artifact failures due to bounded hindsight, bounded foresight, and bounded rationality. Managers' tendencies to overlook anomalies and the organizational and environmental pressures they often experience can result in a rush to judgement to categorize events into pre-defined buckets. Such factors may cause safety drifts, which sow the seeds of artifact failures because such drifts eventually place organizations at the limits, and then beyond

(Oliver, Calvard, & Potočnik, 2017). Focusing attention on the managers' decision processes may aid managers in mitigating biases driving decisions through their technological and governance structure choices. A third promising area for future research would be to examine governance structures that are embedded in organizational systems to identify if managers' decision-making rights are aligned with the requisite expertise to make those decisions, and how a (mis-)match affects the ability of transacting parties to pursue safety goals (Farjoun & Starbuck, 2005). Fourth, future research that builds on theories seeking to explain accidents (Perrow, 1984; Reason, 1997; Weick & Sutcliffe, 2001) may better explain and predict firms' differential learning from artifact failures based on these firms' heterogeneous governance structures. Further, research can also examine the impact of safety culture and changes in regulation on the propensity of high-reliability artifact failure (see, e.g., the forensic accident investigation of Strauch (2002)).

Both single case study (Oliver, *et al.* 2017) and multiple case study (Eisenhardt, 1989; Weick & Sutcliffe, 2001) approaches can be useful in examining high-reliability artifact failures. The investigation of unusual cases of high-reliability artifact failures may be particularly useful in revealing anomalies, new information, and identifying blind-spots of decision-makers (Behfar & Okhuysen, 2018). Research using large datasets (Baum & Dahlin, 2007) can make important contributions by leveraging substantial exogenous shocks (e.g., regulatory changes) in the research design to pinpoint better the causal factors that impact organizational system safety.

3.4. Conclusion

High-reliability artifacts are ubiquitous, and often are embedded in dynamic environments characterized by high technical system complexity (e.g., railroads and space exploration). These artifacts inevitably experience some form of failure. While attempts at organizational adaptation to cope with systems complexity and bounded rationality through the division of labor, organizational boundaries, and governance structures enable adaptation, such attempts to provide organizational solutions will also constrain adaptation as *organizational* complexity rises. Such difficulties arise, in part, because of the evolution over time of inefficient or inappropriate partitioning of decision rights, which impedes adaptive responses to mitigate recurrences of artifact failure.

In addition, managers' responsible for managing high-reliability artifacts are limited by bounds on cognition. On the one hand, managers have *bounded hindsight,* which impedes their ability to discern causal connections between events leading up to artifact failure. On the other hand, managers have *bounded foresight*, which limits their ability to devise safeguards adequately to mitigate artifact failure originating in non-behavioral factors. Managers are also boundedly rational, which increases their likelihood of being unaware of anomalies, as well as a proclivity to fit those anomalies that they do become aware of into their pre-existing mental schemas.

There is also a tendency of managers to posit that suspending operations to uncover the underpinning factors leading to anomalies is often not justified due to the high economic costs the action entails. Managers adaptive decisions are predicated on current (economic) costs of a disturbance and ensuing adaptation, and do not fully consider future private and public costs. Unfortunately, this managerial approach to decision making can fail to mitigate (catastrophic) failures, some of which can be more costly than any adaptive response might have been (e.g., Lac-Megantic). The managers' dilemma on whether to intervene when anomalies occur is exacerbated by the fact that some anomalies and glitches are random events that do not always foreshadow massive failures. The very idea of identifying which anomalies deserve managers' attention is an under-explored area of research in the extant management literature that warrants further inquiry.

This study incorporate theories of accidents and complexity theory into the transaction cost economics approach to mitigate high-reliability artifact failure. Beyond the well-known construct of bounded rationality, I suggest the usefulness of considering the constructs of bounded hindsight and bounded foresight for mitigating artifact failure.

This study suggests ways that managers can learn from past artifact failure to enable the mitigation of recurrences. First, when artifact failure inevitably occurs, my developed framework encourages transacting parties to search for non-behavioral (technological) uncertainty rather than defaulting to attributing the cause to opportunism. Second, the framework proposes that transacting parties eschew the prism of narrow self -interest, and design governance structures that promote their joint interest in achieving safety. Third, the framework recommends that transacting parties give particular attention to anomalies that occur during the operation of complex artifacts, regardless of the immediate economic costs. Fourth, the framework encourages transacting parties to jointly seek (multiple) explanations for anomalies, which can expand their existing schema. Fifth, the framework suggests that organizational systems match (inter-disciplinary) expertise with responsibility and decision rights for contract design. Finally, the framework cautions against the use of complex contracts for the governance of high-reliability artifacts as their usage may impede adaptation. This dissertation study offers a technical systems rationale for artifact failure without considerations of behavioral uncertainty (e.g., opportunism) being invoked.

Given the litigious environment experienced by corporate managers it is understandable that the heuristic of assigning blame has evolved as the dominant logic in response to artifact failure. This dissertation study suggests that there is another way forward. The framework I develop concerning high-reliability artifacts, admittedly has higher implementation cost in the short run relative to the status quo of assigning fault based on an assumption of contractual opportunism and/or gross incompetence by a contractual party. Further, the expected benefits of my suggested change from the status quo cannot be easily quantified because it may not be possible to calculate the economic gains (losses prevented) obtained from safe operations of the artifact. Absent a straightforward cost-benefit analysis, transacting parties may find it arduous to justify and maintain economic costs incurred. Nonetheless, remediableness is possible (Williamson, 1996). Put differently, despite the higher implementation costs, a move from the status quo to the suggested approach championed in the current study can be expected to achieve net gains. There are two primary reasons for why a positive net gain is anticipated. First, high reliability artifact failure can lead to massive economic costs (e.g., Deep Water Horizon) if not bankruptcy. Second, improvements in safety not only lead to the reduced likelihood of artifact failures but also improve operations, which increase revenues.

As our use of complex and high-reliability artifacts increases, managers need to recognize that some failures are unpredictable and inevitable, and originate in the underpinning complex socio-technical system. Nonetheless, improved managerial attention to anomalies, and improvement in the partitioning of decision control rights to those possessing expertise for the relevant problem at hand will enable organizations to reduce the likelihood of artifact failures. Such organizational change to improve success often will require substantially lower cost than the likely benefits in adopting the developed framework provided in this study, thereby satisfying the remediableness criterion. In the evolving science of organization (Barnard 1938; Williamson, 1996), both research and practice concerning high-reliability systems can do better.

CONCLUSION

In my dissertation, in an effort to improve the architecture and governance of complex artifacts, I sought to develop a deeper understanding of anomalies and failures of (high-reliability) artifacts. In doing so, I revisited some assumptions of theories of design and contracting, specifically 'near decomposability' of complex systems (Simon, 1962), bounded rationality, and opportunism (Williamson, 1985). These constructs of 'near decomposability,' bounded rationality, and opportunism, play a vital role in explaining artifact failures. To aid my investigation into the roles of these constructs in the management of complex artifacts, and explaining complex artifact failures I drew on a wide range of theories including theories of contracting (e.g., Williamson, 1996), complexity theory (e.g., Waldrop, 1993), theories of accidents (e.g. Dekker *et al*. 2011; Weick & Sutcliffe, 2001), team theory (Marschak & Radner, 1972), theories of design (e.g., Simon, 1973; 1996), and others. I proposed current approaches to manage complex artifacts and uncover causes for complex artifact failures may be misplaced as these approaches are predicated on empirically unfounded assumptions regarding near decomposability, managerial cognition, and opportunism. The empirical investigations conducted in the course of my dissertation to evaluate these premises reveal near decomposability reveals that complex systems are not necessarily decomposable, earnest transacting parties may disagree on why anomalies and artifact failures occurred, and not all artifact failures are a result of opportunism. Further, the limits on managerial cognition may be quite severe and impede contracting when managers are dealing with complex technical systems that are not, or cannot, be fully understood.

I submit that in managing complex artifacts that inevitably and unpredictably fail, conducting an empirical examination of the underpinning technology to determine boundaries of decomposability is critical. Once boundaries of decomposability are identified, governance structures can be aligned such that the associated failure complexity is low, i.e., the failures of subsystems are not confounded such that the identification of causes of failures remain beyond managerial cognitive ability. I maintain managerial cognitive ability is limited because managers have bounded hindsight, bounded foresight, are boundedly rational, and prone to errors. Because these managers are often dealing with ill-structured problems (Simon, 1973), anomalies and unexpected artifact failure constitute an independent, i.e., no (human) behavioral source of uncertainty. Failure of managers to recognize non-behavioral uncertainty falsely leads managers to assess that they have a far more comprehensive understanding of complex technical systems than they possess, which can lead to avoidable disagreements, and fracture contracts. Critically, managers have a proclivity to invoke, without evidence, opportunistic behavior by their transacting partners to explain failures that remain a mystery. When managers invoke opportunism without evidence, it impedes these managers from learning about the artefact itself. In the case of high-reliability artefacts, this tendency to default to invoking opportunism as an explanation for failures may cause contract termination, and is counter-productive to mitigating recurring artifact failures. To aid in mitigating high-reliability artifact failure, I suggest managers align decision rights with expertise, develop simpler contracts, and direct managerial attention to anomalies and failures of the high-reliability artifact, irrespective of the immediate economic costs of artifact failure.

As the world we inhabit is populated with artifacts that we do not fully comprehend (March and Augier, 2008), (will) inevitable fail (Perrow, 1984), and lie beyond full managerial control, adopting an analytically convenient but reductionist approach for managing them may not be an appropriate approach. To mitigate the likelihood of artifact failure, especially in the cases of high-reliability artifacts whose failure can cause massive damage, prudence demands that the foundations on which the science of design and contracting about these artifacts rests on empirically sounds premises. After conducting his investigation into the *Challenger* disaster, Nobel Laureate Richard Feynman highlighted importance of sound premises in reasoning about technology, and stated, “for a successful technology, reality must take precedence over public relations, for nature can never be fooled.” Feynman’s words likely apply to all complex artefacts – nature cannot be fooled, and the price we might pay for adopting empirically false premises for analytic convenience without recognizing their implications and qualifying our conclusions may well be very high. I hope my dissertation is but a small step in developing alternative approaches to ‘managing’ complex artifacts.

CONCLUSIÓN

En mi disertación, en un esfuerzo por mejorar la arquitectura y el gobierno de los artefactos complejos, procuré desarrollar una comprensión más profunda de las anomalías y fallas de los artefactos (de alta confiabilidad). Al hacerlo, volví a examinar algunas suposiciones de las teorías de diseño y contratación, específicamente "casi descomponibilidad" de sistemas complejos (Simon, 1962), racionalidad limitada y oportunismo (Williamson, 1985). Estas construcciones de "casi descomponibilidad", racionalidad limitada y oportunismo desempeñan un papel vital en la explicación de las fallas de los artefactos. Para ayudar a mi investigación sobre los roles de estos constructos en el manejo de artefactos complejos, y explicando fallas complejas de artefactos, utilicé una amplia gama de teorías que incluyen teorías de contratación (por ejemplo, Williamson, 1996), teoría de la complejidad (por ejemplo, Waldrop, 1993).), teorías de accidentes (por ejemplo, Dekker et al., 2011; Weick y Sutcliffe, 2001), teoría de equipos (Marschak y Radner, 1972), teorías del diseño (por ejemplo, Simon, 1973; 1996) y otras. Propuse enfoques actuales para manejar artefactos complejos y descubrir las causas de fallas de artefactos complejos que pueden estar fuera de lugar, ya que estos enfoques se basan en suposiciones empíricamente infundadas con respecto a la descomposición cercana, la cognición gerencial y el oportunismo. Las investigaciones empíricas llevadas a cabo en el curso de mi disertación para evaluar estas premisas revelan que la descomposición cercana revela que los sistemas complejos no son necesariamente descomponibles, las partes fiables pueden no estar de acuerdo sobre por qué ocurrieron las anomalías y fallas de artefactos, y no todas las fallas de artefactos son resultado de oportunismo Además, los límites en la cognición gerencial pueden ser bastante severos e impedir la contratación cuando los gerentes están lidiando con sistemas técnicos complejos que no son, o no pueden, ser entendidos completamente.

Afirmo que en el manejo de artefactos complejos que inevitablemente y de forma impredecible fallan, realizar un examen empírico de la tecnología de soporte para determinar los límites de la descomposición es crítico. Una vez que se identifican los límites de la descomposición, las estructuras de gobierno pueden alinearse de modo que la complejidad de la falla asociada sea baja, es decir, las fallas de los subsistemas no se confundan de manera tal que la identificación de las causas de las fallas permanezca más allá de la capacidad cognitiva de gestión. Mantengo que la capacidad cognitiva de gestión es limitada porque los gerentes tienen una visión retrospectiva limitada, una previsión limitada, son racionalmente limitados y están sujetos a errores. Debido a que estos gerentes a menudo se enfrentan a problemas mal estructurados (Simon, 1973), las anomalías y el fallo inesperado de los artefactos constituyen una fuente de incertidumbre independiente (es decir, sin comportamiento). El hecho de que los gerentes no reconozcan la incertidumbre no conductual conduce falsamente a los gerentes a evaluar que tienen una comprensión mucho más completa de los sistemas técnicos complejos de los que poseen, lo que puede conducir a desacuerdos evitables y fracturas en los contratos. Críticamente, los gerentes tienen la tendencia de invocar, sin evidencia, un comportamiento oportunista por parte de sus socios en la transacción para explicar las fallas que siguen siendo un misterio. Cuando los gerentes invocan el oportunismo sin evidencia, esto les impide aprender sobre el propio artefacto. En el caso de los artefactos de alta confiabilidad, esta tendencia a no recurrir al oportunismo como una explicación de los fracasos puede causar la terminación del contrato, y es contraproducente para mitigar los fallos recurrentes de artefactos. Para ayudar a mitigar la falla de artefactos de alta confiabilidad, sugiero a los gerentes alinear los derechos de decisión con la experiencia, desarrollar contratos más simples y dirigir la atención de la gerencia a las anomalías y fallas del artefacto de alta confiabilidad, independientemente de los costos económicos inmediatos de la falla de artefactos.

A medida que el mundo que habitamos está poblado por artefactos que no comprendemos completamente (marzo y agosto de 2008), (será) inevitablemente fallará (Perrow, 1984) y quedará fuera del control administrativo total, adoptando un enfoque analíticamente conveniente pero reduccionista para la gestión. Puede que no sean un enfoque apropiado. Para mitigar la probabilidad de fallas en los artefactos, especialmente en los casos de artefactos de alta confiabilidad cuya falla puede causar un daño masivo, la prudencia exige que los fundamentos sobre los cuales la ciencia del diseño y la contratación de estos artefactos se basen en premisas empíricamente sonoras. Después de realizar su investigación sobre el desastre del Challenger, el premio Nobel Richard Feynman destacó la importancia de las premisas sólidas en el razonamiento sobre la tecnología y declaró: "para una tecnología exitosa, la realidad debe tener prioridad sobre las relaciones públicas, porque la naturaleza nunca puede ser engañada". Palabras de Feynman probablemente se apliquen a todos los artefactos complejos: la naturaleza no puede ser engañada, y el precio que podríamos pagar por adoptar premisas empíricamente falsas para la conveniencia analítica sin reconocer sus implicaciones y calificar nuestras conclusiones puede ser muy alto. Espero que mi tesis no sea más que un pequeño paso en el desarrollo de enfoques alternativos para "administrar" artefactos complejos.

REFERENCES

Agrawal, A., Muthulingam, S., & Rajapakshe, T. 2017. How sourcing of interdependent components affects quality in automotive supply chains. ***Production and Operations Management***, 26(8): 1512-1533.

Alchian, A.A., & Demsetz, H. 1972. Production, information costs, and economic organization. ***American Economic Review***, *62*(5): 777-795.

Alchian, A.A., & Woodward, S. 1988. The firm is dead; long live the firm a review of Oliver E. Williamson's the economic institutions of capitalism. ***Journal of Economic Literature***, *26*(1): 65-79.

Amundson, S.D. 1998. Relationships between theory-driven empirical research in operations management and other disciplines. ***Journal of Operations Management***, *16*(4): 341-359.

Anand, G., & Gray, J.V. 2017. Strategy and organization research in operations management. ***Journal of Operations Management***, 53-56: 1-8.

Argyres, N., & Bigelow, L. 2007. Does transaction misalignment matter for firm survival at all stages of the industry life cycle? ***Management Science***, 53(8): 1332-1344.

Argyres, N.S., & Liebeskind, J.P. 1999. Contractual commitments, bargaining power, and governance inseparability: Incorporating history into transaction cost theory. ***Academy of Management Review***, *24*(1): 49-63.

Argyres, N.S., & Mayer, K.J. 2007. Contract design as a firm capability: An integration of learning and transaction cost perspectives. ***Academy of Management Review***, *32*(4): 1060-1077.

Arrow, K.J. 1969. The organization of economic activity: issues pertinent to the choice of market versus nonmarket allocation. ***The Analysis and Evaluation of Public Expenditure: The PPB System***, *1*: 59-73.

Arrow, K.J. 1974. ***The Limits of Organization***. New York: WW Norton & Company.

Arrow, K.J. 1985. The economics of agency. In J. Pratt and R.J. Zeckhauser (Eds.), ***Principals and Agents*****:** ***The Structure of Business:*** 37-51. Boston: Harvard Business School Press.

Augier, M., &. March, J.G. 2008. Realism and comprehension in economics: A footnote to an exchange between Oliver E. Williamson and Herbert A. Simon. ***Journal of Economic Behavior & Organization***, *66*(1): 95-105.

Balakrishnan, S., & Wernerfelt, B. 1986. Technical change, competition and vertical integration. ***Strategic Management Journal***, *7*(4): 347-359.

Baldwin, C.Y., & Clark, K.B. 2000. ***Design Rules: Volume 1. The Power of Modularity***. Cambridge, MA: MIT Press.

Baum, J.A., & Dahlin, K.B. 2007. Aspiration performance and railroads' patterns of learning from train wrecks and crashes. ***Organization Science***, *18*(3): 368-385.

Behfar, K., & Okhuysen, G.A. 2018. Perspective—Discovery within validation logic: Deliberately surfacing, complementing, and substituting abductive reasoning in hypothetico-deductive inquiry. ***Organization Science***, *29*(2): 323-340.

Bhardwaj, A., & Ketokivi, M. 2016. How decomposable are nearly decomposable systems? The case of locomotive engine failures. In J. Humphreys (Ed.), ***Proceedings of the Seventy-sixth Annual Meeting of the Academy of Management***, Online ISSN: 2151-6561.

Boardman, M.E. 2005. Contra Proferentem: The Allure of Ambiguous Boilerplate. ***Michigan Law Review***, *104*, 1105.

Calvano, C.N., & John, P. 2004. Systems engineering in an age of complexity. ***Systems Engineering***, 7(1): 25-34.

Campbell, B. 2018. ***The Lac-Mégantic Rail Disaster: Public Betrayal, Justice Denied***. Toronto: James Lorimer & Company.

Cao, Z., & Lumineau, F. 2015. Revisiting the interplay between contractual and relational governance: A qualitative and meta-analytic investigation. ***Journal of Operations Management***, *33*(1): 15-42.

Chen, I.J., & Paulraj, A. 2004. Towards a theory of supply chain management: the constructs and measurements. ***Journal of Operations Management***, *22*(2): 119-150.

Chikudate, N. 2009. If human errors are assumed as crimes in a safety culture: A lifeworld analysis of a rail crash. ***Human Relations***, *62*(9): 1267-1287.

Choi, T.Y., & Krause, D.R. 2006. The supply base and its complexity: Implications for transaction costs, risks, responsiveness, and innovation. ***Journal of Operations Management***, 24(5): 637-652.

Choi, T.Y., & Krause, D.R. 2006. The supply base and its complexity: Implications for transaction costs, risks, responsiveness, and innovation. ***Journal of Operations Management***, *24*(5): 637-652.

Clearfield, C., & Tilcsik, A. 2018. ***Meltdown: Why Our Systems Fail and What We Can Do About It.*** New York: Penguin Press.

Coase, R.H. 1937. The nature of the firm. ***Economica***, *4*(16): 386-405.

Coase, R.H. 1994. ***Essays on Economics and Economists***. Chicago: University of Chicago Press.

Coase, R.H. 2006. The conduct of economics: The example of Fisher Body and General Motors. ***Journal of Economics & Management Strategy***, *15*(2): 255-278.

Colfer, L.J., & Baldwin, C.Y. 2016. The mirroring hypothesis: Theory, evidence, and exceptions. ***Industrial and Corporate Change***, 25: 709-738.

Conklin, T. 2016. ***Pre-accident Investigations: An Introduction to Organizational Safety***. New York: CRC Press.

Cox, D.R. 1972. Regression models and life tables. ***Journal of the Royal Statistical Society, Series B***, 34(2): 187-220.

Crook, T.R., Combs, J.G., Ketchen, D.J., & Aguinis, H. (2013). Organizing around transaction costs: What have we learned and where do we go from here? ***Academy of Management Perspectives****, 27*(1): 63-79.

Dekker, S. 2011. ***Drift into Failure: From Hunting Broken Components to Understanding Complex Systems***. Farnham, UK: Ashgate.

Dekker, S., Cilliers, P., & Hofmeyr, J.-H. 2011. The complexity of failure: Implications of complexity theory for safety investigations. ***Safety Science***, 49(6): 939-945.

Dhanorkar, S.S., Siemsen, E., & Linderman, K.W. 2017. Promoting change from the outside: Directing managerial attention in the implementation of environmental improvements. ***Management Science,*** *64*(6): 2535-2566.

Dul, J., Hak, T., Goertz, G., & Voss, C. 2010. Necessary condition hypotheses in operations management**.** ***International Journal of Operations & Production Management***, *30*(11): 1170-1190.

Dunbar, R.L.M., & Starbuck, W.H. 2006. Learning to design organizations and learning from designing them. ***Organization Science***, 17(2): 171-178.

Eggleston, K., Posner, E.A., & Zeckhauser, R.J. 2000. ***Simplicity and Complexity in Contracts***. Working Paper No. 93. John M. Olin Program in Law and Economics, Stanford, CT.

Eisenhardt, K.M. 1989. Building theories from case study research. ***Academy of Management Review***, *14*(4): 532-550.

Eisenhardt, K.M., and Graebner E.M. 2007. Theory building from cases: Opportunities and challenges. ***Academy of Management Journal***, 50(1): 25-32.

Ellram, L.M., Tate, W.L., & Billington, C. 2008. Offshore outsourcing of professional services: A transaction cost economics perspective. ***Journal of Operations Management***, *26*(2): 148-163.

Emery, G.W., & Marques, M.A. 2011. The effect of transaction costs, payment terms and power on the level of raw materials inventories. ***Journal of Operations Management***, *29*(3): 236-249.

Farjoun, M. 2005. Organizational learning and action in the midst of safety drift: revisiting the space shuttle program's recent history. In W.H. Starbuck and M. Farjoun (Eds.). ***Organization at the Limit: Lessons from the Columbia Disaster*:** 20-40. Oxford, UK: Blackwell Publishing.

Feynman, R.P. 1986. Appendix F: Personal observations on the reliability of the shuttle. ***Report of the Presidential Commission on the Space Shuttle Challenger Accident***, *2*: F1-F5.

Foss, N.J. 2003. Bounded rationality in the economics of organization: Much cited and little used. ***Journal of Economic Psychology***, *24*(2): 245-264.

Foss, N.J., & Weber, L. 2016. Moving opportunism to the back seat: Bounded rationality, costly conflict, and hierarchical forms. ***Academy of Management Review***, *41*(1): 61-79.

Freeman, R.E., Harrison, J.S., Wicks, A.C., Parmar, B.L., & De Colle, S. (2010). ***Stakeholder Theory: The State of the Art***. Cambridge: Cambridge University Press.

Furubotn, E.G., & Richter, R. 2000. ***Institutions and Economic theory: The Contribution of the New Institutional Economics***. Ann Arbor. MI: University of Michigan Press.

Garicano, L., & Rayo, L. 2016. Why organizations fail: models and cases. ***Journal of Economic Literature***, *54*(1): 137-92.

Gray, J.V., & Handley, S.M. 2015. Managing contract manufacturer quality in the presence of performance ambiguity. ***Journal of Operations Management***, *38*(1): 41-55.

Groop, J., Ketokivi, M., Gupta, M., & Holmström, J. 2017. Improving home care: Knowledge creation through engagement and design. ***Journal of Operations Management***, 53-56: 9-22.

Grover, V., & Malhotra, M.K. 2003. Transaction cost framework in operations and supply chain management research: theory and measurement. ***Journal of Operations Management***, 21(4): 457-473.

Guajardo, J. A., Cohen, M.A., & Netessine, S. 2016. Service competition and product quality in the U.S. automobile industry. ***Management Science***, 62(7): 1860-1877.

Gupta, S., & Krishnan, V. 1999. Integrated component and supplier selection for a product family. ***Production and Operations Management***, 8(2): 163-182.

Haller, C.L., Coward, R. N., & LaMastra, D. 1994. Root cause investigation of cracked cylinder liners in ALCO 251 diesel engines in nuclear standby service. ***SAE Transactions***, 103(3): 1899-1907.

Handley, S.M. 2017. How governance misalignment and outsourcing capability impact performance. ***Production and Operations Management***, *26*(1): 134-155.

Handley, S.M., & Benton, W.C. 2009. Unlocking the business outsourcing process model. ***Journal of Operations Management***, *27*(5): 344-361.

Handley, S.M., & Benton W.C. 2012. The influence of exchange hazards and power on opportunism in outsourcing relationships. ***Journal of Operations Management***, *30*(1-2): 55-68.

Hayek, F.A. 1937. Economics and knowledge. ***Economica***, *4*(13): 33-54.

Hayek, F.A. 1945. The use of knowledge in society. ***American Economic Review***, *35*(4): 519-530.

Hayek, F.A. 1964. In M. Bunge (Eds.). ***The Critical Approach to Science and Philosophy:*** 332-349. London, UK: Transaction Publishers.

Hayek, F.A. 1979. ***The Counter-revolution of Science: Studies on the Abuse of Reason***. Indianapolis, IN: Liberty Fund.

Hempel, C.G. 1965. ***Aspects of Scientific Explanation and Other Essays in the Philosophy of Science***. New York: Free Press.

Hodgson, G.M. 2004. Opportunism is not the only reason why firms exist: Why an explanatory emphasis on opportunism may mislead management strategy. ***Industrial and Corporate Change***, *13*(2): 401-418.

Holcomb, T.R., & Hitt, M.A. (2007). Toward a model of strategic outsourcing. ***Journal of Operations Management***, *25*(2): 464-481.

Hollenbeck, J.R., & Wright, P.M. 2017. Harking, sharking, and tharking: Making the case for post hoc analysis of scientific data. ***Journal of Management***, *43*(1): 5-18.

Hollnagel, E. 1993. ***Human Reliability: Analysis, Context and Control***. London, UK: Academic Press.

Hoopes, D.G., & Postrel, S. 1999. Shared knowledge, “glitches,” and product development performance. ***Strategic Management Journal***, *20*(9): 837-865.

Hosmer, D. W., & Lemeshow, S. 1999. ***Applied Survival Analysis: Regression Modeling of Time to Event Data***. New York: Wiley.

Jacobs, M.A., & Swink, M. 2011. Product portfolio architectural complexity and operational performance: Incorporating the roles of learning and fixed assets. ***Journal of Operations Management***, 29(7-8): 677-691.

Janis, I. 1991. Groupthink. In E. Griffin (Ed.), ***A First Look at Communication Theory***: 235-246. New York: McGraw-Hill.

Jap, S.D., & Anderson, E. 2003. Safeguarding interorganizational performance and continuity under ex post opportunism. ***Management Science***, *49*(12): 1684-1701.

Kang, M.P., Mahoney, J.T., & Tan, D. 2009. Why firms make unilateral investments specific to other firms: The case of OEM suppliers. ***Strategic Management Journal,*** *30*(2): 117-135.

Ketchen, D.J., & Hult, G.T.M. 2007. Bridging organization theory and supply chain management: The case of best value supply chains. ***Journal of Operations Management***, *25*(2): 573-580.

Ketokivi, M., & Choi, T. 2014. Renaissance of case research as a scientific method. ***Journal of Operations Management***, *32*(5): 232-240.

Ketokivi, M., & Mahoney, J.T. 2016. Transaction cost economics as a constructive stakeholder theory. ***Academy of Management Learning & Education***, *15*(1): 123-138.

Ketokivi, M., Mantere, S., & Cornelissen, J. 2017. Reasoning by analogy and the progress of theory. ***Academy of Management Review***, 42(4): 637-658.

Kistruck, G.M., Morris, S.S., Webb, J.W., & Stevens, C.E. 2015. The importance of client heterogeneity in predicting make-or-buy decisions. ***Journal of Operations Management***, *33*(1): 97-110.

Klein, B., Crawford, R.G., & Alchian, A.A. 1978. Vertical integration, appropriable rents, and the competitive contracting process. ***Journal of Law and Economics***, *21*(2): 297-326.

Klein, P.G., Mahoney, J.T., McGahan, A.M., & Pitelis, C.N. 2019. Organizational governance adaptation: Who is in, who is out, and who gets what. ***Academy of Management Review***, *44*(1): 6-27.

Knight, F.H. 1921. ***Risk, Uncertainty and Profit***. New York: Hart, Schaffner and Marx.

Kreps, D.M. 1996. Corporate culture and economic theory. In J.P. Buckle and J. Michie (Eds.). ***Firms, Organizations and Contracts:*** 221-275. Oxford, UK: Oxford University Press.

Lampel, J., & Giachetti, C. 2013. International diversification of manufacturing operations: Performance implications and moderating forces. ***Journal of Operations Management***, *31*(4): 213-227.

Leibenstein, H. 1976. ***Beyond Economic Man.*** Cambridge, MA: Harvard University Press.

Leveson, N. 2004. A new accident model for engineering safer systems. ***Safety Science***, 42: 237-270.

Leveson, N. 2004. A new accident model for engineering safer systems. ***Safety Science***, *42*(4): 237-270.

Liu, Y., Luo, Y., & Liu, T. 2009. Governing buyer–supplier relationships through transactional and relational mechanisms: Evidence from China. ***Journal of Operations Management***, *27*(4): 294-309.

Llewellyn, K.N. 1931.What price contract? An essay in perspective. ***Yale Law Journal***, *40*(5): 704-751.

Lonsdale, C. 2001. Locked-in to supplier dominance: On the dangers of asset specificity for the outsourcing decision. ***Journal of Supply Chain Management***, *37*(1): 22-27.

Macaulay, S. 1963. Non-contractual relations in business: a preliminary study. ***American Sociological Review***, 28(1): 55-67.

MacCormack, A., & Mishra, A. 2015. Managing the performance trade-offs from partner integration: Implications of contract choice in R&D Projects. ***Production and Operations Management***, *24*(10): 1552-1569.

MacDuffie, J.P. 1997. The road to "root cause": Shop-floor problem-solving at three auto assembly plants. ***Management Science***, 43(4): 479-502.

Macneil, I. 1974. The many futures of contract. ***Southern California Law Review***, *47*(2): 688-816.

Mahapatra, S.K., Narasimhan, R., & Barbieri, P. 2010. Strategic interdependence, governance effectiveness and supplier performance: A dyadic case study investigation and theory development. ***Journal of Operations Management***, *28*(6): 537-552.

Mahoney, J.T. 2005. ***Economic Foundations of Strategy***. Thousand Oaks, CA: Sage Publications.

Mantel, S.P., Tatikonda, M.V., & Liao, Y. 2006. A behavioral study of supply manager decision-making: Factors influencing make versus buy evaluation. ***Journal of Operations Management***, *24*(6): 822-838.

Mantere, S., & Ketokivi, M. 2013. Reasoning in organization science. ***Academy of Management Review***, *38*(1): 70-89.

March, J.G. 1994. ***Primer on Decision Making: How Decisions Happen***. New York: Free Press.

March, J.G., Sproull, L.S., & Tamuz, M. 1991. Learning from samples of one or fewer. ***Organization Science***, *2*(1): 1-13.

Marschak, J., & Radner, R. 1972. ***Economic Theory of Teams***. New Haven, CT: Yale University Press.

Mayer, K.J., & Argyres. N.S. 2004. Learning to contract: Evidence from the personal computer industry. ***Organization Science***, *15*(4): 394-410.

McIvor, R. 2009. How the transaction cost and resource-based theories of the firm inform outsourcing evaluation. ***Journal of Operations Management***, *27*(1): 45-63.

McLachlin, R. 1997. Management initiatives and just-in-time manufacturing. ***Journal of Operations Management***, *15*(4): 271-292.

McMillan, C.J., & Overall, J.S. 2017. Crossing the chasm and over the abyss: Perspectives on organizational failure. ***Academy of Management Perspectives***, *31*(4): 271-287.

Melnyk, S.A., Pagell, M., Jorae, G., & Sharpe, A.S. 1995. Applying survival analysis to operations management: Analyzing the differences in donor classes in the blood donation process. ***Journal of Operations Management***, 13(4): 339-356.

Mentzer, J.T., DeWitt, W., Keebler, J.S., Min, S., Nix, N.W., Smith, C.D., & Zacharia, Z.G. 2001. Defining supply chain management. ***Journal of Business Logistics***, *22*(2): 1-25

Merkert, R., Smith, A.S., & Nash, C.A. 2012. The measurement of transaction costs—evidence from European railways. ***Journal of Transport Economics and Policy***, *46*(3): 349-365.

Milgrom, P., & Roberts, J. 1987. Informational asymmetries, strategic behavior, and industrial organization. ***American Economic Review***, *77*(2): 184-193.

Mill, J.S. 1882. ***A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation***. New York: Harper & Brothers.

Monczka, R.M., Petersen, K.J., Handfield, R.B., & Ragatz, G.L. 1998. Success factors in strategic supplier alliances: the buying company perspective. ***Decision Sciences***, *29*(3): 553-577.

Monteverde, K., & Teece, D. J. 1982. Supplier switching costs and vertical integration in the automobile industry. ***Bell Journal of Economics***, 13(1): 206-213.

Nickerson, J.A., & Zenger, T.R. 2004. A knowledge-based theory of the firm—The problem-solving perspective. ***Organization Science***, *15*(6): 617-632.

Nickerson, R.S. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. ***Review of General Psychology***, *2*(2): 175-220.

Novak, S., & Eppinger, S. 2001. Sourcing by design: Product complexity and the supply chain. ***Management Science***, 47(1): 189-204.

Novak, S., & Stern, S. 2008. How does outsourcing affect performance dynamics? Evidence from the automobile industry. ***Management Science***, 54(12): 1963-1979.

Novak, S., & Wernerfelt, B. 2012. On the grouping of tasks into firms: Make-or-buy with interdependent parts. ***Journal of Economics & Management Strategy***, 21(1): 53-77.

Oliver, N., Calvard, T., & Potočnik, K. 2017. Cognition, technology, and organizational limits: Lessons from the Air France 447 disaster. ***Organization Science***, 28(4): 729-743.

Park, J.-K., & Ro, Y. K. 2011. The impact of a firm's make, pseudo-make, or buy strategy on product performance. ***Journal of Operations Management***, 29(4): 289-304.

Peerally, M. F., Carr, S., Waring, J., & Dixon-Woods, M. 2017. The problem with root cause analysis. ***BMJ Quality & Safety***, 26: 417-422.

Perrow, C. 1984. ***Normal Accidents***. Princeton, N.J.: Princeton University Press.

Polanyi, M. 1958. ***Personal Knowledge, Towards a Post Critical Epistemology***. Chicago, IL: University of Chicago.

Poppo, L., & Zenger, T. 2002. Do formal contracts and relational governance function as substitutes or complements? ***Strategic Management Journal***, *23*(8): 707-725.

Ramasesh, R.V., & Browning, T.R. 2014. A conceptual framework for tackling knowable unknown unknowns in project management. ***Journal of Operations Management***, *32*(4): 190-204.

Ramdas, K., & Randall, T. 2008. Does component sharing help or hurt reliability? An empirical study in the automotive industry. ***Management Science***, *54*(5): 922-938.

Reason, J. 1997. ***Managing the Risks of Organizational Accidents***. London, UK: Routledge.

Rozeboom, W.W. 1997. Good science is abductive, not hypothetico-deductive. In L. L. Harlow, S. A. Mulaik, & J. H. Steiger (Eds.), ***What If There Were No Significance Tests?***: 335-391. Mahwah, N.J.: Lawrence Erlbaum Associates.

Rubin, P.H. 1993. ***Managing Business Transactions***. New York: Free Press.

Salvador, F., Chandrasekaran, A., & Sohail, T. 2014. Product configuration, ambidexterity and firm performance in the context of industrial equipment manufacturing. ***Journal of Operations Management***, 32(4): 138-153.

Sanchez, R., & Mahoney, J.T. 1996. Modularity, flexibility, and knowledge management in product and organization design. ***Strategic Management Journal***, 17(Winter special issue): 63-76.

Sarkis, J., Zhu, Q., & Lai, K.H. 2011. An organizational theoretic review of green supply chain management literature. ***International Journal of Production Economics***, *130*(1): 1-15.

Schelling, T.C. 1960. ***The Strategy of Conflict***. Cambridge, MA: Harvard University Press.

Simon, H.A. 1957. ***Models of Man; Social and Rational***. Oxford, UK: Wiley

Simon, H.A. 1962. The architecture of complexity. ***Proceedings of the American Philosophical Society***, 106(6): 467-482.

Simon, H.A. 1973. The structure of ill structured problems. ***Artificial Intelligence***, *4*(3): 181-201.

Simon, H.A. 1976. ***Administrative Behavior: A Study of Decision-making Processes in Administrative Organization*** (3rd ed.). New York: The Free Press.

Simon, H.A. 1985. Human nature in politics: The dialogue of psychology with political science. ***American Political Science Review***, *79*(2): 293-304.

Simon, H.A. 1996. ***The Sciences of the Artificial*** (3rd ed.). Cambridge, MA: MIT Press.

Simon, H.A. 2000. Bounded rationality in social science: Today and tomorrow. ***Mind & Society***, *1*(1): 25-39.

Smith, A. 1776. ***An Inquiry into the Wealth of Nations***. London, UK: Strahan and Cadell.

Sosa, M.E., Mihm, J., & Browning, T.R. 2013. Linking cyclicality and product quality. ***Manufacturing & Service Operations Management***, *15*(3): 473-491.

Starbuck, W., & Farjoun, M. (Eds.). 2005. ***Organization at the Limit: Lessons from the Columbia Disaster***. Oxford, UK: Blackwell Publishing.

Starbuck, W.H., & Milliken, F.J. (1988). Challenger: fine-tuning the odds until something breaks. *Journal of Management Studies*, *25*(4): 319-340.

Sterman, J.D. 1994. Learning in and about complex systems. ***System Dynamics Review***, *10*(2): 291-330.

Strauch, B. 2002. ***Investigating Human Error: Incidents, Accidents, and Complex Systems***. New York: Ashgate Publishing.

Terpend, R., & Krause, D.R. 2015. Competition or cooperation? Promoting supplier performance with incentives under varying conditions of dependence**. *Journal of Supply Chain Management***, *51*(4): 29-53.

Therrien, M.C., Valiquette-L'Heureux, A., Normandin, J.M., & Smits, P. 2016. Tightly coupled governance for loosely coupled wicked problems: the train explosion in Lac-Mégantic case. ***International Journal of Risk Assessment and Management***, *19*(4): 260-277.

Thompson, J.D. 1967. ***Organizations in Action: Social Science Bases of Administrative Theory***. New York: McGraw Hill.

Toulmin, S.E. 2003. ***The Uses of Argument*** (updated ed.). Cambridge, MA: Cambridge University Press.

Tucker, A.L. 2007. An empirical study of system improvement by frontline employees in hospital units. ***Manufacturing & Service Operations Management***, 9: 492-505.

Tucker, A.L., & Edmondson, A.C. 2003. Why hospitals don't learn from failures: Organizational and psychological dynamics that inhibit system change. ***California Management Review***, 45: 55-72.

Tucker, A.L., & Edmondson, A.C. 2003. Why hospitals don't learn from failures: Organizational and psychological dynamics that inhibit system change. ***California Management Review***, *45*(2): 55-72.

Ülkü, S., & Schmidt, G.M. 2011. Matching product architecture and supply chain configuration. ***Production and Operations Management***, *20*(1): 16-31.

Ulrich, K. 1995. The role of product architecture in the manufacturing firm. ***Research Policy***, 24(3): 419-440.

Ulrich, K.T., & Ellison, D.J. 2005. Beyond make-buy: Internalization and integration of design and production. ***Production and Operations Management***, 14(3): 315-330.

Van de Ven, A.H. 2007. ***Engaged Scholarship: A Guide for Organizational and Social Research***. Oxford, UK: Oxford University Press.

Vaughan, D. 2005. System effects: On slippery slopes, repeating negative patterns, and learning from mistake? In W.H. Starbuck and M. Farjoun (Eds.). ***Organization at the Limit: Lessons from the Columbia Disaster***: 41-59. Oxford, UK: Blackwell Publishing.

Villena, V.H., & Craighead, C.W. 2017. On the same page? How asymmetric buyer–supplier relationships affect opportunism and performance. ***Production and Operations Management***, *26*(3): 491-508.

Wagner, S.M., & Bode, C. 2014. Supplier relationship-specific investments and the role of safeguards for supplier innovation sharing. ***Journal of Operations Management***, 32(3): 65-78.

Waldrop, M M. 1993. ***Complexity: The Emerging Science at the Edge of Order and Chaos***. New York: Simon and Schuster.

Wang, M., Zhang, Q., Wang, Y., & Sheng, S. (2016). Governing local supplier opportunism in China: Moderating role of institutional forces. ***Journal of Operations Management***, *46*(1): 84-94.

Wason, P.C. 1960. On the failure to eliminate hypotheses in a conceptual task. ***Quarterly Journal of Experimental Psychology***, *12*(3): 129–40.

Weick, K.E. 1989. Theory construction as disciplined imagination. ***Academy of Management Review***, *14*(4): 516-531.

Weick, K.E. 2005. Making sense of blurred images: Mindful organizing in mission STS-107. In W.H. Starbuck and M. Farjoun (Eds.). ***Organization at the Limit: Lessons from the Columbia Disaster***: 159-177. Oxford, UK: Blackwell Publishing.

Weick, K.E., & Sutcliffe, K.M. 2001. ***Managing the Unexpected: Resilient Performance in an Age of Uncertainty***. San Francisco, CA: John Wiley & Sons.

Williamson, O.E. 1971. The vertical integration of production: Market failure considerations. ***American Economic Review***, 61(2): 112-123.

Williamson, O.E. 1975. ***Markets and Hierarchies***. New York: The Free Press.

Williamson, O.E. 1985. ***The Economic Institutions of Capitalism***. New York: Free Press.

Williamson, O.E. 1991. Comparative economic organization: The analysis of discrete structural alternatives. ***Administrative Science Quarterly***, *36*(2): 269-296.

Williamson, O.E. 1993a. Calculativeness, trust, and economic organization. ***Journal of Law and Economics***, *36*(1): 453-486.

Williamson, O.E. 1993b. Opportunism and its critics. ***Managerial and Decision Economics***, *14*(2): 97-107.

Williamson, O.E. 1996. ***The Mechanisms of Governance.*** Oxford, UK: Oxford University Press.

Williamson, O.E. 1998. Transaction cost economics: How it works; where it is headed. ***De Economist***, *146*(1): 23-58.

Williamson, O.E. 2008a. Outsourcing: Transaction cost economics and supply chain management. ***Journal of Supply Chain Management***, *44*(2): 5-16.

Williamson, O.E. 2008b. Transaction cost economics. In Menard. C., & Shirley. M.M. (Eds.), ***Handbook of New Institutional Economics***: 41-65. Berlin, Heidelberg: Springer.

Williamson, O.E., Wachter, M.L., & Harris, J.E. 1975. Understanding the employment relation: The analysis of idiosyncratic exchange. ***Bell Journal of Economics***, *6*(1): 250-278.

Wilson, P.F., Dell, L.D., & Anderson, G.F. 1993. ***Root Cause Analysis: A Tool for Total Quality Management***. Milwaukee, WI: ASQ Press.

Wolf, C. 1993. ***Markets or Governments: Choosing between Imperfect Alternatives***. Cambridge, MA: MIT Press.

Wu, A.W., Lipshutz, A.K.M., & Pronovost, P.J. 2008. Effectiveness and efficiency of root cause analysis in medicine. ***JAMA***, 299(6): 685-687.

Zeckhauser, R.J., & Pratt, J.W. (Eds.) 1985. ***Principals and Agents: The Structure of Business***. Boston: Harvard Business School Press.

Zipkin, P. 2012. A reply to Williamson's "outsourcing…". ***Production and Operations Management***, *21*(3): 465-469.

Zu, X., & Kaynak, H. 2012. An agency theory perspective on supply chain quality management. ***International Journal of Operations & Production Management***, *32*(4): 423-446.