14-05

September 10, 2014

ESMT Working Paper

CONTRACTING IN MEDICAL EQUIPMENT MAINTENANCE SERVICES

AN EMPIRICAL INVESTIGATION

TIAN CHAN, INSEAD
FRANCIS DE VÉRICOURT, ESMT
OMAR BESBES, COLUMBIA UNIVERSITY

ISSN 1866-3494

Abstract

Contracting in medical equipment maintenance services: An empirical investigation

Author(s):* Tian Chan, INSEAD
Francis de Véricourt, ESMT
Omar Besbes, Columbia University

Equipment manufacturers offer different types of maintenance service plans (MSPs) that delineate payment structures between equipment operators and maintenance service providers. These MSPs allocate risks differently and thus induce different kinds of incentives. A fundamental question, therefore, is how such structures impact service performance and the service chain value. We answer empirically this question. Our study is based on a unique panel data covering the sales and service records of over 700 diagnostic medical body scanners. By exploiting the presence of a standard warranty period, we overcome the key challenge of isolating the incentive effects of MSPs on service performance from the confounding effects of adverse selection. We found that moving an operator from a basic pay-per-service plan to a fixed-fee full-protection plan leads to both a reduction in reliability and an increase in service costs. We further show that the increase in cost is driven by both the operator and the service provider. Our results point to the presence of losses in service chain value in the maintenance of medical equipment, and provide the first evidence that a basic pay-per-service plan, where the risk of equipment failure is borne by the operator, can actually improve performance and costs.

Keywords: Maintenance repair, service contracting, co-production, empirical operations management, service chain value, healthcare industry

* Contact: Francis de Véricourt, ESMT, Schlossplatz 1, 10178 Berlin,
Phone: +49 30 21231-1517, francis.devericourt@esmt.org.

1. Introduction

Operators of capital-intensive equipment often devote a large annual budget on maintenance so as to ensure high equipment reliability. For instance, in the medical imaging equipment industry, which forms the basis of our study, a top-of-the-line Computed Tomography (CT) or Magnetic Resonance Imaging (MRI) equipment typically costs US$1 million each, and require annual maintenance expenses corresponding to 10% of their price (ECRI Institute 2013). Many leading manufacturers, such as General Electric Co., Siemens AG, and Hewlett-Packard Co., have therefore expanded their maintenance service offerings over the last decade to maintain revenue levels as competition in the core product market intensifies (Sawhney et al. 2004).

A key feature of maintenance services is that value is co-produced, i.e., both proper operational handling by the operator and proper execution of maintenance routines by the service provider have an effect on reliability–and hence collaboration between the equipment operator and the service provider is critical to create good service outcomes (typically measured through failure rate and maintenance costs). This collaboration, in turn, crucially depends on the maintenance service plan (MSP) that contractually determines the payment structure for the service and hence may influence both parties' behavior in their interactions with the equipment. Indeed, MSPs take many forms, ranging from a pay-per-service contract where a markup for labor and parts is charged every time a repair service is required (basic MSP), to a fixed-price full-protection plan covering all services over an agreed period (full-protection MSP). Under the basic MSP, because the operator pays each time a failure occurs, she bears a greater risk of equipment failure. In contrast, under the full-protection MSP, she only pays an upfront fixed-fee, and it is the service provider who now carries the risk of equipment failure.

This paper provides one of the first empirical analyses of the impact of MSPs on maintenance service outcomes. By exploiting a unique dataset of a major medical device manufacturer, we disentangle the incentive effects induced by MSPs from issues of endogenous contract selection. This enables us to quantify the (relative) effect of these MSPs on failure rate, onsite visits, remote resolutions as well as labor and replacement part costs. We find that a pay-per-service plan significantly improves these operational performance metrics over a fixed-price full-protection plan. In turn, this implies that pay-per-service maintenance plans outperform full-protection plans at the service chain level (i.e., treating the service provider and the operator as a single economic unit). To the best of our knowledge, this is the first study that empirically establishes such a phenomenon. Through mediation analysis, we further establish that both the service provider and operator contributes to the improvement in service costs.

Despite the key role MSPs play in the co-production of maintenance services, almost no empirical study have explored the relative performance of these contracts. There are, indeed, good reasons for this lack of empirical studies to date. First, datasets of MSPs amenable to analysis are rare and lacking. Beside the obvious issue of confidentiality, these datasets are fairly recent, as manufacturers have only recently started to foray into maintenance services in a significant way (Sawhney et al. 2004). In fact, manufacturers often face challenges in managing the transition into service provision and thus are generally slow and cautious in expanding their after-sales service offerings (Oliva and Kallenberg 2003). Second, maintenance contracts can lead to both adverse-selection and moral hazard. Indeed, operators with higher usage demand and hence failure rate, for instance, may tend to select a full-protection plan over a basic one. Similarly, because both parties' actions affect the chance of failure, there is a double moral hazard problem: both the operator and service provider may have incentives to shirk and not properly care for the equipment. Therefore, a significant challenge lies in developing identification strategies that simultaneously account for adverse selection and (double-sided) moral hazard (Chiappori and Salanié 2003, Abbring et al. 2003).

We have been able to overcome these challenges by obtaining a rich dataset covering the monthly service records of more than 700 medical imaging equipment (MRI and CT Scanners) in 441 different hospitals from a major medical device manufacturer. Hospitals operate the equipment, while the manufacturer provides maintenance services. This panel data records each new sale of equipment, which comes with a standard one-year warranty period. At the end of this period, operators elect to sign on to one of the three types of MSPs (basic, partial and full-protection plans), each one of them providing the hospital and manufacturer with different incentives.

In terms of the fundamental identification challenge above, we exploit the structure of this unique dataset to account for adverse selection in the estimation of the incentive effects of MSPs. Since during the warranty period, all hospitals have the same incentive structure, any observed differences in service failures and costs among the hospitals over this period of time can be attributed to differences in their innate operating conditions. Based on this, we adapt a fixed effects model to estimate the innate operating conditions of each piece of equipment, and therefore remove its confounding effects in our final estimates. This enables us to measure the impact of contract type on service outcomes due to (double-sided) moral hazard, i.e., the effect on service outcomes due to the shifting of equipment failure risk from one party to the other.

We find that moving an operator from a pay-per-service basic MSP to a full-protection MSP results in a significant increase (30%) of operator reported failure rate. Further, the service provider increases by 80% the number of onsite visits to the operator on full-protection MSP, with associated increases of 69% in the total time spent onsite and 102% in the cost of materials and spares.

We further show that, given a reported failure, the service provider makes 46% more onsite visits (as opposed to remotely solving the problem) to an operator on full-protection MSP compared to an operator on basic MSP. Additionally, all the increase in labor and spares costs can be explained by the increase in onsite visits, meaning that conditional on the service provider visiting, the expenditure on labor and materials appears independent on the contract type.

Finally, using a smaller sample of data in which operator usage is available, we show that MSPs have no significant effect on equipment usage. This suggests that the increase in failures observed (going from the basic MSP to the full-protection MSP) is not driven by the operator increasing equipment usage, but by other actions of the operator that affects failure rates, e.g., exerting less care in equipment handling.

Overall, we find evidence that MSPs induce changes in behavior on *both* the operator and service provider. On the part of the operator, we find evidence of operator moral hazard, so that the full-protection MSP increases the number of operator reported failures. In addition, we exhibit evidence of service provider moral hazard, in that we see an increase in likelihood that the service provider will respond to a given failure by dispatching an engineer to make an onsite visit to the operator on full-protection. These effects in turn lead to a substantial increase in the level of labor and material expenditures.

Our study appears to be the first to establish that the basic MSP can be beneficial for maintenance outcomes and costs. This contradicts the prevailing wisdom, both by scholars and practitioners, that manufacturers intending to move into services should assume more of the equipment failure risks (Oliva and Kallenberg 2003, Deloitte Research 2006). Thus, they should offer the full-protection MSP, which allows the manufacturer to assume responsibility for the cost of equipment repairs, or even adding performance-based provisions to it (i.e., service level guarantees or payments tied with performance measures such as failure rate or uptime), which allows the manufacturer to additionally assume the responsibility for service performance. This advice has found some empirical support. Specifically in the aviation industry, Guajardo et al. 2012 found that aircraft engine units on performance-based contracts have higher reliability (i.e., lower failure rates) compared to those on basic contracts. Our results, however, provide a caution against this prevailing wisdom, as we observed both lower reliability and higher service costs when the service provider assumes higher responsibilities over equipment failures. Because we did not observe any significant change in usage, from a service chain value perspective, this implies that the full-protection MSP leads to a lower value compared to the basic MSP.

Finally, Guajardo et al. 2012 theorized that the increases in reliability they observed empirically is due to the service provider exerting costly effort (e.g., spending a longer time during preventive maintenance routines). Thus, their work hypothesizes a tradeoff (i.e., a negative relationship) between performance and cost across contracts. Our work shows that a positive relationship can exist, i.e., contracts that produce better reliability at a cheaper cost can co-exist with contracts that are worse off in both measures.

The rest of the paper is structured in the following manner. First, we discuss in §2 the relevant literature and theory in contracts and maintenance service provisions. We then discuss the industrial setting and data in §3. Next, we present a graphical representation of the results in §4, which then leads to the empirical approach employed in §5. We show our main results in §6 and a set of additional results / robustness checks in §7. Finally, we conclude and discuss some implications in §8.

2. Literature Review

Most theoretical research on contracting in the operations management literature focuses on the context of physical supply chains and manufacturing systems (see, e.g., Cachon 2003, Nagarajan and Sošić 2008 for reviews of work in this area) and there is comparatively less theoretical work on service contracting (Zhou and Ren 2011).

Within the service literature, our work falls in the area of maintenance service contracting. Some papers have analyzed maintenance contracting problems in a principal-agent framework, where the operator (customer) is the principal and the service provider the agent. In particular, Plambeck and Zenios (2000) studies a dynamic principal-agent model where the principal (the customer) outsources maintenance work of an equipment to a service provider who repairs the equipment when it breaks down. The customer is assumed to have no influence over equipment failures (and thus is free from moral hazard issues), but the service provider is susceptible to moral hazard because his effort in performing repairs influences maintenance outcomes, and cannot be observed. Similarly, Kim et al. (2007) uses a principal-agent framework (in a single interaction setup) to study a maintenance setting where a customer contracts with multiple service providers who can influence maintenance outcomes by deciding on the inventory level of spares parts.

In general, however, services are co-produced in the sense that both the operator and the service provider contribute to value creation (Fuchs and Leveson 1968). From a contracting perspective, this can lead to double moral hazard issues when efforts of both agents are not contractible. Jain et al. (2013), for example, consider a setup such as this to study the issue of payments uncertainty when the service provider is small scaled. In this stream of research, the theoretical work of Roels et al. (2010) is perhaps closest to our paper. Indeed, Roels et al. (2010) develop a double moral hazard model to compare the service chain value between a basic, full-protection, and performance-based plan. The key finding is that, depending on the context, each one of these three kinds of contracts can theoretically dominate the other two. The performance-based plan dominates when concerns of moral hazard from both parties are important. In contrast, the basic plan (resp. full-protection plan) dominates when concerns of moral hazard comes mainly from the operator (resp. service provider). Our work provides some of the first empirical evidence for these theoretical predictions.

If the theoretical literature on maintenance service contract is scarce, empirical work in operations management on the topic is almost nonexistent despite the practical importance of these issues in industry. To the best of our knowledge, Guajardo et al. (2012) is the only paper that empirically explores the performance of maintenance service contracts. The setup of our respective works differ along two dimensions. First, the contracts analyzed are of a different nature. Guajardo et al. (2012) focus on performance-based contracts (versus basic plans), which are common in aerospace and defense industries, where equipment owners are typically very large entities. In contrast, we focus on full-protection plans (versus basic plans), which are much more common in the medical equipment industry, or other industries in which equipment owners are smaller. Second, Guajardo et al. (2012) focus on estimating the effects of MSPs on only one metric of service performance, the failure rate; we analyze failure rate as well, but also study service costs (onsite visits, labor hours, and spares cost) and usage rates (the latter on a smaller set of data). This helps us identify changes in both party's behavior over not just service performance but also on costs and usage. The conclusions of the two papers are also very different. Guajardo et al. (2012) find in the aircraft engine industry that performance-based contracts improve reliability over time-and-material plans that are similar to our basic MSP. In other words, they find that basic MSPs hamper performance, which they attribute to service provider moral hazard, i.e., the service provider taking less care on maintenance routines for operators on basic plans. In contrast, we find in our set-up that a basic MSP improves both performance and cost, driven by both the service provider and operator's behavior.

More generally, scholars have also empirically studied a variety of issues pertaining to supply chain coordination. For example, empirical papers have examined the effects of product component sharing (Ramdas and Randall 2008), information sharing (Terwiesch et al. 2005, Cui et al. 2013), or vertical integration (Novak and Stern 2008, 2009) on collaboration outcomes. While we share a common goal of identifying configurations that optimize collaboration outcomes, our paper focuses particularly on contractual incentives.

Finally, the economics literature has a number of papers that explore the association between service contracts and performance (see Chiappori and Salanié 2003 for a review of works in the area). As Chiappori and Salanié (2003) point out, similar to the operations management literature, the number of empirical papers lag behind their theoretical counterparts. Particularly, papers on service contracts appear to focus more on project-based services–e.g. IT software development (Banerjee and Duflo 2000, Gopal et al. 2003, Kalnins and Mayer 2004, Susarla et al. 2010, Susarla 2012), and offshore drilling (Corts and Singh 2004). The one-off relationship framed around a project that characterizes project-based services differs significantly from the dynamics of repeated interactions found in the maintenance industry.

3. The After-Sales Maintenance of Medical Equipment

We obtained MSP and service performance data for a major manufacturer's CT and MRI scanner business for a major OECD country. As abovementioned, CT and MRI scanners are complex pieces of equipment, with top of the line models costing typically in excess of US$1 million. These pieces of equipment have a typical lifespan of around ten years; during this time regular quarterly preventive maintenance visits by a service provider (who is the manufacturer in this case) are required to ensure smooth operations. Operators (radiologists of hospitals) also play a role in equipment reliability through proper equipment handling.[1]

Each new equipment sale comes with a standard warranty period typically of one year, at the end of which operators elect to sign on to one of the three types of MSPs. The three types of MSPs, in decreasing order of coverage, are (i) the full-protection MSP, offering the operator complete protection against both spares and labor costs incurred by product failure in return for a fixed annual fee, (ii) a partial-protection MSP, in which the operator pays a fixed fee of labor but has to pay a variable amount for any spares consumed during the repair process, and (iii) the basic MSP, which is preventive maintenance only, and requires operators to pay for both labor and spares incurred for repairs; this is a time-and-materials contract with a low fixed-fee. Table 1 highlights the key differences across the three types.

Table 1 Summary of MSP Coverage

	Preventive Maintenance	Labor Charges	Material Charges
Full-Protection	Yes	Yes	Yes
Partial-Protection	Yes	Yes	
Basic	Yes		

Figure 1 depicts the high-level process of how equipment failures are reported and resolved. The process begins with an operator calling in to the manufacturer's service call center to report failure events. Electronic tickets are created at this point that log equipment details and a description of the problem. The manufacturer has a team of service call engineers, who then call the operators to understand the problem. Combined with remote access to the equipment to diagnose and resolve problems, a large fraction of tickets can be resolved at this stage. For problems that cannot be resolved, the ticket is then assigned to a service engineer with the suitable skillset who then travels to the operator premises to

[1] For example, MRI surface coils, a small but mission-critical component for localized anatomy scanning, is a delicate component and requires particular care in use (Kawaja and Durmis 2005).

resolve the problem. After the problem is resolved, the ticket is closed with information on the total repair hours and material cost (spares and/or consumables) used during repairs logged into the system.

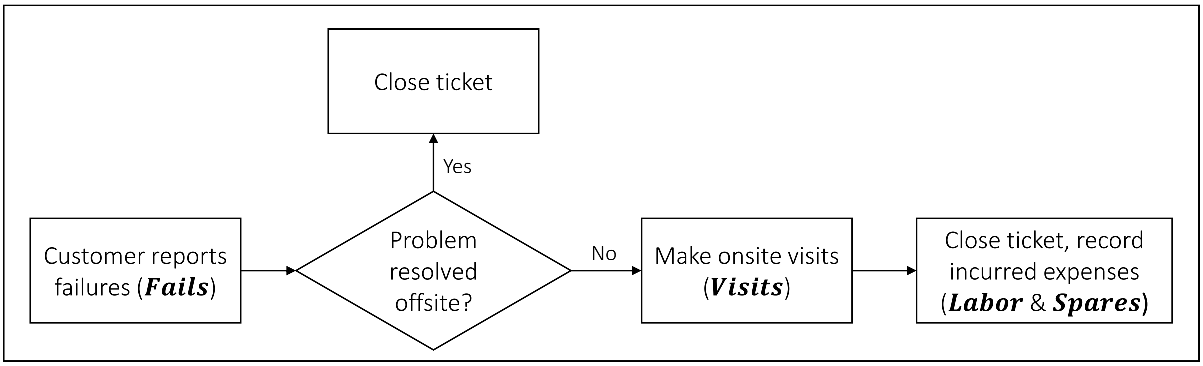

Figure 1 High-level Process of Equipment Failure Reporting and Resolution.

3.1. Data

Our data covers MSP and service performance / operating cost data for the manufacturer's CT and MRI scanners sold between 2008 and 2012 for a major OECD country. The data covers 712 pieces of equipment, of which 57% are CT scanners. Among the operators, 74% select the full-protection MSP after the warranty period, 21% select the basic MSP, and the remaining 5% choose the intermediate partial-protection MSP. Because our data also cover the standard warranty period for each equipment, at any point in time a piece of equipment can be under one of four contract types, which we code as dummy variables ($Full, Partial, Basic$, and $Warranty$ respectively).

We track four measures for service performance and operating costs on a monthly basis, which cover all the key elements of the process represented in Figure 1. The first measure is $Fails$, a count variable measuring the number of reported failures. The second measure $Visits$ captures the number of onsite visits made to the operator premises to resolve problems (some problems may require more than one visit, while some problems can be resolved offsite without any visits). The variables $Labor$ and $Spares$ capture the operating cost of problem resolution–$Labor$ captures the total labor hours spent onsite, while $Spares$ captures the cost of consumables and spares for maintenance and repairs[2].

[2] We scaled *Labor* and *Spares* with a fixed positive multiplicative factor to mask the absolute magnitude of the results. This transformation does not have an effect on our estimation at all given our use of an exponential setup which reports results as percentage changes from a given base.

Table 2 provides summary statistics for the variables. The service performance and operating cost variables $Fails, Visits, Labor$, and $Spares$ are both non-negative (minimum achieved at zero), and have a large right-skew (signified by a maximum that is typically many times larger than the mean). For example, the average number of *Labor* hours consumed in a month for a piece of equipment is fairly low, at 2.8 hours. However, there are occasions when a large amount of *Labor* hours need to be expended–the observed maximum is 219 hours, which is 78 times the average. Non-negative right-skewed dependent variables are frequently encountered in econometric studies and we will discuss in §5 a standard approach that we adopt to model them.

In any observed month, a piece of equipment must fall into exactly one of the four types of service plans $Warranty, Full, Partial$ or *Basic*. Thus, the four variables are mutually exclusive as exactly one variable can take the value of one in any observed month of any piece of equipment. The mean of $Warranty$–0.44–indicates that 44% of the months in our data set consist of observations of the equipment when it is covered under standard warranty. Because the standard warranty is a period where all operators have to go through, we observe more months for equipment covered under warranty than under $Full$ (43%), $Partial$ (2%), or $Basic$ (11%) plans.

Finally, we also keep track of the *Age* of the equipment in months starting from its installation date, which is also the date when the standard warranty period begins.

Table 3 shows rank correlation statistics for the variables. Note that the service performance and operating cost variables $Fails, Visits, Labor$, and $Spares$ are positively rank correlated. So months with a higher number of reported failures will tend to have a higher number of onsite visits and a higher number of labor and spares expended.

4. An Initial Visual Representation of the Main Results

Our goal is to identify the impact of MSPs on service performance and operating costs. Figure 2 traces the average number of failures (on a quarterly basis) as equipment crosses the warranty period into the MSP period for three groups of operators. Three lines are plotted: the solid line corresponds to operators who selected the *Full* contract, the dashed line to operators who selected the *Partial* contract, and the dotted line to operators who chose the *Basic* contract. The shaded period represents the period where equipment is still under standard warranty.

Table 2 Summary Statistics. 712 equipment. 23,173 monthly observations.

Variables	Description	Mean (S.D.)	Min-Max
Service Performance and Operating Cost			
Fails	Number of failures reported per month	1.43 (1.79)	0 - 18
Visits	Number of site visits per month	1.00 (2.04)	0 - 45
Labor	Hours consumed per month	2.80 (8.21)	0 - 219
Spares	Cost of spares consumed per month	1,502 (7,631)	0 - 173,930
Service Contract (Plan)			
Warranty	Indicator–operator is on standard warranty	0.44 (0.50)	0-1
Full	Indicator–operator is on full plan	0.43 (0.49)	0-1
Partial	Indicator–operator is on partial plan	0.02 (0.15)	0-1
Basic	Indicator–operator is on basic plan	0.11 (0.32)	0-1
Equipment Characteristic			
Age	Age of equipment in months	19.1 (13.0)	1-60

Table 3 Spearman Rank Correlation Coefficient. 712 pieces of equipment. 23,173 monthly observations.

	Fails	*Visits*	*Labor*	*Spares*	*Warranty*	*Full*	*Partial*	*Basic*	*Age*
Fails	1.00								
Visits	0.75	1.00							
Labor	0.71	0.94	1.00						
Spares	0.53	0.67	0.63	1.00					
Warranty	0.13	0.11	0.12	0.07	1.00				
Full	-0.03	-0.01	-0.02	0.02	-0.76	1.00			
Partial	0.00	0.00	0.00	-0.03	-0.13	-0.13	1.00		
Basic	-0.16	-0.16	-0.15	-0.12	-0.32	-0.31	-0.05	1.00	
Age	-0.15	-0.12	-0.13	-0.05	-0.73	0.58	0.10	0.18	1.00

We observe that the failure rate tends to decrease over time across the three groups of operators. This general trend suggests that *non-repeating problems*, such as failures that are reported due to operator unfamiliarity, or equipment defects during manufacturing, are gradually discovered and resolved over time, leading to a gradual reduction in failure rate. Thus, our empirical model will need to control for age effects on service outcomes.

Figure 2 further highlights two key aspects of the data. First, when we focus purely on the data in the warranty period (say, months -3 to -1 in the shaded area), we see that the average number of operator reported failures is much higher for operators who select *Full* (1.60) compared to operators who select *Basic* (1.10). The difference (of 0.50 failures per month) between the two groups of operators during the warranty period provides visual evidence of adverse contract selection, i.e., operators do not choose MSPs randomly but

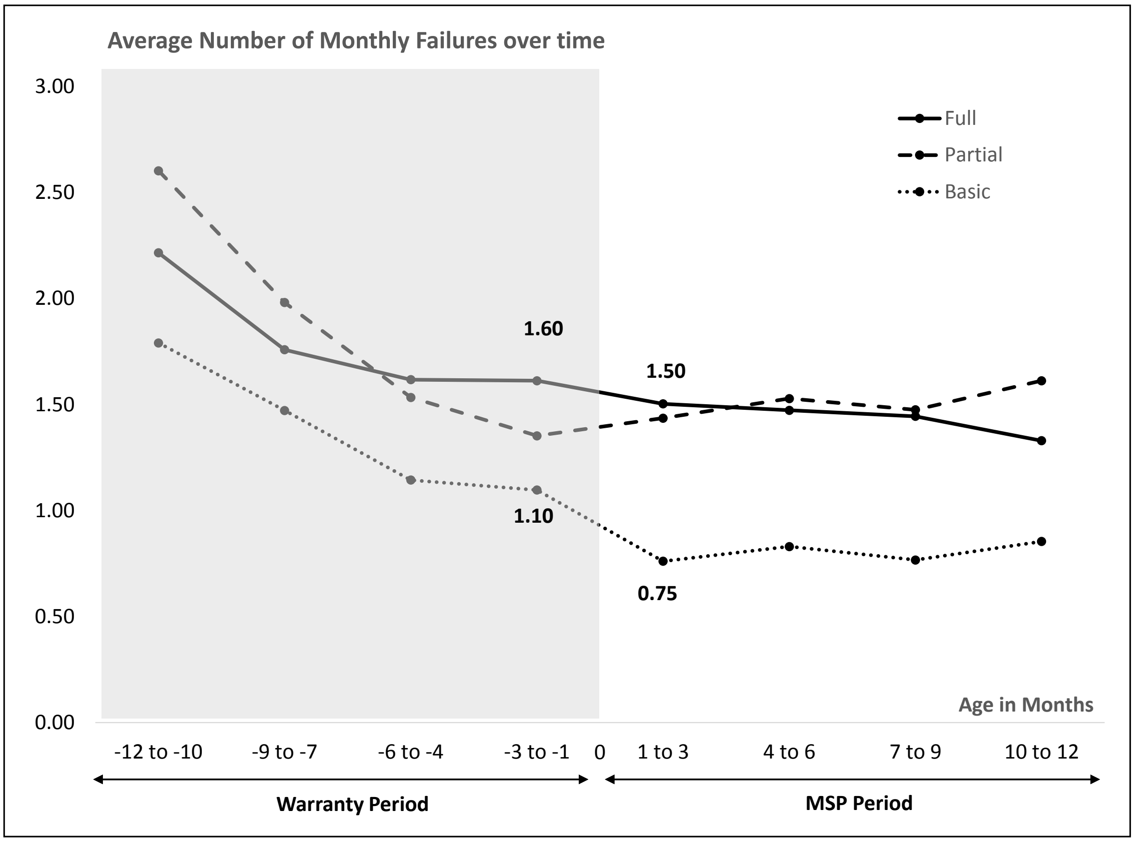

Figure 2 **A plot of average quarterly number of failures over time (time 0 is the time when warranty expires).**

instead choose MSPs with more protection if they have innately higher failure rates (and thus benefit more from it).

Second, if we focus on the dotted line representing equipment failure rates for customers on *Basic*, we can visually observe a sharp drop in reported failures as the operators transition from the warranty period into the MSP period (from 1.10 in months '-3 to -1' to 0.75 in months '1 to 3'). This drop is less apparent for customers who select *Full* (1.60 to 1.50). These concurrent changes results in a widening of the gap between the lines representing customers who select *Full* as opposed to customers who select *Basic*, as customers transition form the warranty to the MSP period. This widening of the gap (e.g., from $1.60 - 1.10 = 0.50$ to $1.50 - 0.75 = 0.75$) is indicative of the effects of MSPs on service performance due to the different incentive structure of the MSPs. The rigorous measurement of effects such as these is the main question we will be after in the subsequent sections.

5. Estimation approach

In this section we present a model that incorporates the possibility of endogenous MSP selection, which leads to the identification strategy we use to estimate the impact of MSPs on different service metrics.

For each piece of equipment i, one can think of the expected costs of the equipment as depending on three key variables: the current $Plan$ the equipment is on, its Age, and an individual constant c_i measuring the innate propensity for equipment failure (which is in turn dependent on equipment usage characteristics). The variable $Plan$ influences service outcomes because of incentives and thus is central to the specification. As seen in Figure 2, Age appears to be a strong determinant of equipment failure rate. Finally, every piece of equipment may have differing operating characteristics (such as load), and we incorporate a time-independent constant for each piece of equipment. We denote the expected service outcome C of equipment i observed in a given month as $E[C|Plan, Age, i]$ (where C can represent service performance measure $Fails$ or service cost measures $Visits, Labor$, and $Spares$).

For the purpose of discussion, assume for now that $Full$ and $Basic$ are the only MSP plans offered to the operator. The central goal is to measure the effects of MSP on the expected service costs when hypothetically moving an operator from $Basic$ to $Full$. Denote this quantity as ΔC_B^F. Given the potential of double moral hazard, ΔC_B^F is influenced by the actions of both the operator and the service provider. Thus, one can think of ΔC_B^F as being driven by two components: $\Delta C_B^F = \Delta C_c + \Delta C_s$. The first is the effect of operator actions, ΔC_c (one would expect that ΔC_c to be non-negative because the operator would rationally exert effort to reduce the costs she bears), and the second is the effect of service provider actions ΔC_s (one would expect that ΔC_s to be non-positive). Our work aims to estimate ΔC_B^F, i.e., the net effect of MSP on service cost due to the actions of both the service provider and the operator.

With the goal of estimation properly defined, we can thus express $E[C|Plan, Age, i]$ as follows.

$$\begin{aligned} E[C|Full, Age, i] &= c_i + \Delta C_B^F + \beta_t Age \\ E[C|Basic, Age, j] &= c_j + \beta_t Age \end{aligned} \quad (1)$$

A problem of endogeneity arises since individual operator characteristics c_i's are unobserved. If c_i is correlated with $Plan$, then a simple comparison of service costs across operators on $Full$ against operators on $Basic$ (controlling for Age) will be biased. To illustrate, suppose there are only two types of operators, one (denoted i) who has a large innate demand of scans, and another one (denoted by j) with a low demand. Because operator i is expected to incur much higher service costs compared to operator j ($c_i > c_j$), she finds it more economical to purchase the $Full$ MSP. Here, a simple comparison between the service cost between operators on $Full$ and operators on $Basic$, taking account of equipment age, would lead to a biased estimate of the impact of MSP on service metrics, i.e. $\Delta C_B^F + c_i - c_j$ instead of ΔC_B^F.

Our empirical approach critically relies on being able to observe a standard warranty period. In such a period, the service provider serves a population of operators that *face the same incentives.* While moral hazard issues may arise during this period, the effect on the service metric C, denoted as ΔC_w would be the same. Hence, in this period, we have that for equipment i and j:

$$\begin{aligned} E[C|Warranty, Age, i] &= c_i + \Delta C_w + \beta_t Age \\ E[C|Warranty, Age, j] &= c_j + \Delta C_w + \beta_t Age. \end{aligned} \tag{2}$$

The essential insight is that the warranty period offers a way to account for operational characteristics of each individual operator c_i. Figure 3 depicts the elements defined above graphically. The impact of moving from the *Basic* to the *Full* contract (ΔC_B^F) is expressed by a change in the size of the gap as operators cross over to the MSP period. Specifically, $\Delta C_B^F > 0$ indicates dominant operator moral hazard while $\Delta C_B^F < 0$ implies dominant service provider moral hazard.

Define I_w and I_f as indicator variables denoting if the operator is at the moment on $Warranty$ or $Full$ plan respectively. Equations (1, 2) can be summarized by Equation 3 below

$$E[C|Plan, Age, i] = c_i + \Delta C_w I_w + \Delta C_B^F I_f + \beta_t Age \tag{3}$$

Note that the equation takes the form of a fixed effects equation. Essentially, we isolate ΔC_B^F from the innate differences in failure propensity $c_i - c_j$ due to adverse selection by estimating the unobserved innate characteristics of each operator through a fixed-effect (c_i

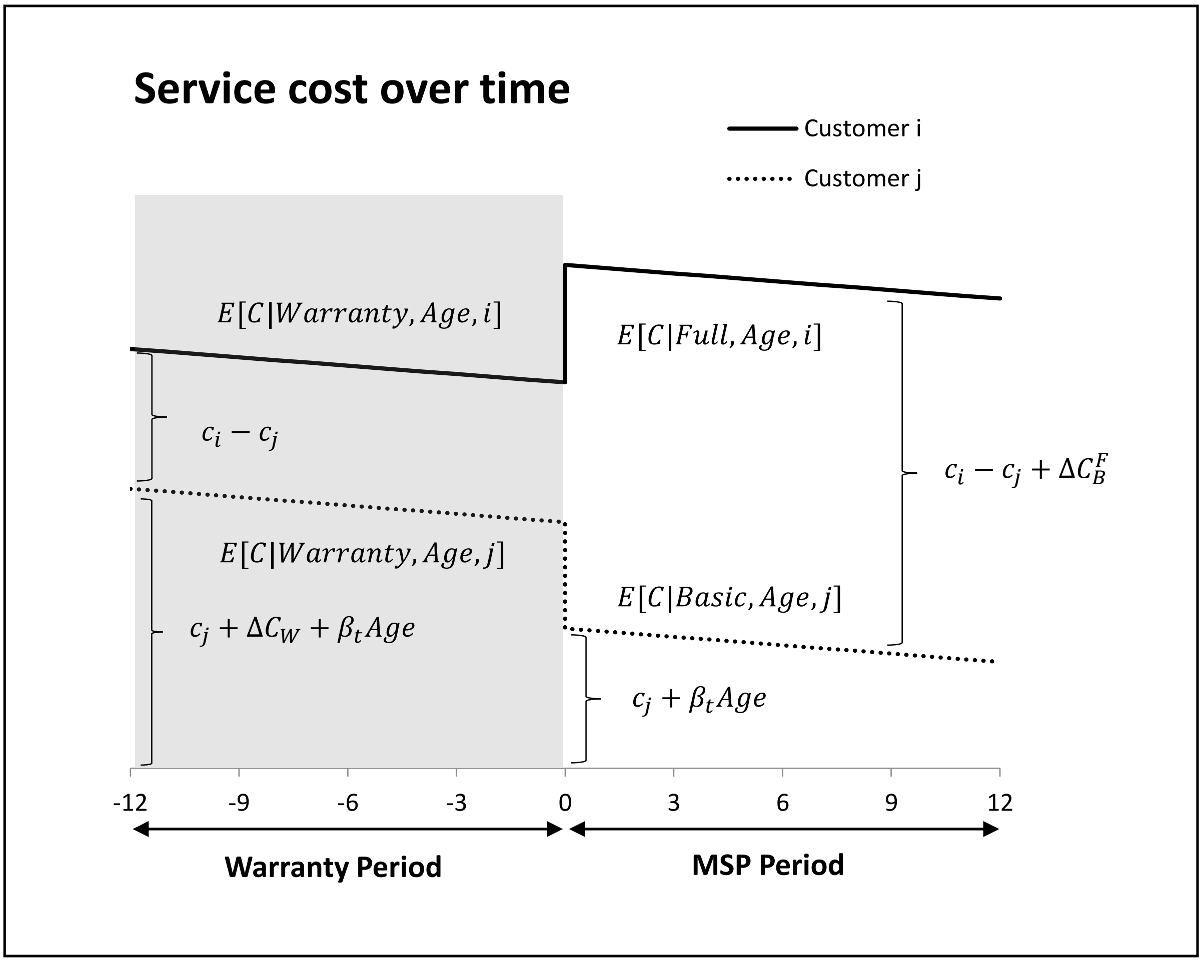

Figure 3 **A plot of service costs over time with endogenous selection and double moral hazard.**

for each operator i), controlling for a common time-trend applied to all operators and the current service plan of each operator.

Note also that the coefficient of interest, ΔC_B^F, could alternatively be obtained via a difference-in-difference specification (i.e., the difference between the two operators during the MSP period $c_i - c_j + \Delta C_B^F$ minus the difference during the warranty period $c_i - c_j$). Our choice of the fixed effect specification over difference-in-difference stems from the availability of *individual panels*. Using fixed effects allows us to model the significant heterogeneity across operators *on the same MSP* (in §7.1, we show visually that customers on the same MSP can have very different usage patterns).

There are two key assumptions to the model. First, c_i is assumed to be time-invariant. Using a smaller sample of data where equipment usage is observed over time, we show in §7.1 that usage patterns, a key determinant of failure rate, is stable across time. Second, we assume that the effect of moral hazard on outcomes is similar for every customer. Specifically, there should no higher-order interactions between c_i with ΔC_w and ΔC_B^F. This assumption prevents inclusion of a large number of interaction variables and is essential to maintain a parsimonous econometric model.

Our data also contains a *Partial* plan. Let I_p denote an indicator for operators who are currently on the *Partial* plan and let ΔC_B^P be defined as the impact on the service metric C of moving a operator from *Basic* to *Partial*. We are now ready to introduce the full specification of the model we estimate.

$$E[C|Plan, Age, i] = \exp[c_i + \Delta C_w I_w + \Delta C_B^F I_f + \Delta C_B^P I_p + \beta_t Age]. \tag{4}$$

We use Generalized Linear Model (GLM) with a log link (i.e. enclose the right-hand-side with an exponent) as a basis for estimation.[3] We also assume that the dependent variable, conditional on the independent variables, has a Poisson distribution. Estimation based on Poisson distributions has many desirable robustness properties. First, the estimator is consistent even if the Poisson distributional assumption does not hold (Wooldridge 2010). In addition, because our model uses *fixed effects*, a potential problem of *incidental parameters* arises–meaning that approaches that use dummy variables (unconditional likelihood estimation) may not converge to the correct solution. This is an issue particular to short panels in non-linear models. The Poisson estimator does not suffer from this problem (Cameron and Trivedi 2009). In all cases we report cluster-robust errors, which produce standard errors that are also robust to the potential presence of heteroskedasticity and auto-correlation.

6. Results

6.1. The Impact of MSP on Service Outcomes

Table 4 presents the results for the regressions. Consider first the model with dependent variable *Fails*. To focus our discussion, let us look first at the coefficients for *Full*, which is the effect on service outcomes of moving an operator from *Basic* to *Full*. The analysis shows a statistically significant ($p < 0.001$) and positive effect on the number of reported failures of being on the full-protection contract, compared to the basic contract.

[3] GLM generalizes the familiar least-squares model and encompasses many non-linear estimators, e.g. Logistic, Poisson and Negative Binomial. Especially, Poisson and Negative Binomial both uses log links. Particularly, GLM with log-links are well-suited to handle non-negative right-skewed dependent variables.

Table 4 Effect of MSP on service performance and operational costs

DV:	$Fails$	$Visits$	$Labor$	$Spares$
Model:	GLM with log link and Poisson distribution			
Warranty	$0.40(0.05)^{***}$	$0.72(0.09)^{***}$	$0.63(0.13)^{***}$	$0.60(0.22)^{**}$
Full (ΔC_B^F)	$0.26(0.05)^{***}$	$0.59(0.09)^{***}$	$0.52(0.13)^{***}$	$0.70(0.23)^{**}$
Partial (ΔC_B^P)	$0.21(0.07)^{**}$	$0.39(0.14)^{**}$	$0.46(0.19)^{*}$	$-0.08(0.46)^{\mathrm{ns}}$
Basic	BASELINE			
Age	$-0.01(0.00)^{***}$	$-0.01(0.00)^{***}$	$-0.02(0.00)^{***}$	$-0.00(0.00)^{\mathrm{ns}}$
Fixed Effects	Yes	Yes	Yes	Yes
Observations	23,173	23,115	23,115	22,473
Log-likelihood	-33,240	-31,802	-92,230	-71,119,603

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, $^{\mathrm{ns}}$ not significant

Standard Errors are clustered by equipment

This implies dominant operator moral hazard. Furthermore, the estimate of ΔC_B^F, 0.26 corresponds to an increase in failure rate of 30%.[4]

Consider now the model with dependent variable $Visits$. Here, we also see a significant ($p < 0.001$) effect on the number of site visits made by the service provider on operators with $Full$ (as opposed to $Basic$). In addition, for a 30% increase in failure rate, the service provider makes an additional 80% increase in onsite visits. That the effect size for $Visits$ is greater than $Fails$ suggest that given a reported failure, the service provider is more likely to visit an operator on $Full$ (we test this more formally in the next section).

If we focus our attention on the coefficients for $Full$ on the *Labor* and $Spares$ models, we see statistically significant increases in both *Labor*($p < 0.001$) and *Spares*($p < 0.01$). Specifically, *Labor* increases by 69%, while $Spares$ increases by 102%.

Overall, our analysis establishes that operator moral hazard is dominant. While it is not possible to fully disentangle operator from service provider moral hazard, the results also suggest the presence of non-trivial service provider moral hazard, which is apparently not driven by cost concerns. We reiterate the point that the service provider is more likely to treat a given failure by making an onsite visit (as opposed to resolving the problem remotely) for operators on the $Full$ plan. This leads to corresponding increases in labor and spares outlay. Interestingly, for operators on $Partial$ plan (where the service provider

[4] The translation from a coefficient β to effect size in an exponential model is given by $\exp[\beta] - 1$.

bears the cost of labor hours but not spares), we see only significant increases in labor hour expenditure. So, consistent across all plans is the observation that the service provider spends more when he bears the portion of the cost.

One question is whether the increase in failure rate observed is due to increases in equipment usage. Using a smaller dataset covering 471 observed months over 22 pieces of equipment, we show in §7.1 that operators do not seem to exhibit significant moral hazard behavior in terms of usage–if any exists, the effect is likely to be very small, 2%(-13–18%), and certainly cannot explain the large increases in failure rates we found here. This result is consistent with the empirical results of Ning et al. (2014), which similarly found no moral hazard effects on usage in pay-per-print services.

Thus, given the deterioration in service performance and increase in service costs, and a lack of positive impact on equipment usage, our results establish that the full-protection plan results in an unambiguous decrease in service chain value relative to the basic (time-and-material) plan.

6.2. The Role of Onsite Visits: A Mediation Analysis

The previous analysis suggested that (i) the service provider is performing more $Visits$ per reported failure, and (ii) the main driver of increases in *Labor* and *Spares* appears to be the increase in $Visits$. In this section we deploy mediation analysis (Baron and Kenny 1986) to test these claims. Essentially, by controlling for $Fails$ in the regression for $Visits$, we investigate if increased failure reports fully explain the increase in $Visits$, and by controlling for $Visits$ in the regressions for *Labor* and *Spares* we see if increased $Visits$ fully explain the increases in *Labor* and *Spares*.

The modified specifications is as follows.

$$\begin{aligned}
\log E[Visits|Plan, Age, i] =& c_i + \beta_w I_w + \beta_f I_f + \beta_p I_p + \beta_t Age + \beta_l LogFails \\
\log E[Labor|Plan, Age, i] =& c_i + \beta_w I_w + \beta_f I_f + \beta_p I_p + \beta_t Age + \beta_l LogVisits \\
\log E[Spares|Plan, Age, i] =& c_i + \beta_w I_w + \beta_f I_f + \beta_p I_p + \beta_t Age + \beta_l LogVisits.
\end{aligned} \tag{5}$$

On the right hand side of the regression, we define $LogFails := Log(Fails + 1)$ and $LogVisits := Log(Visits + 1)$ given that many months have zero failures or visits.

Table 5 reports the results of the analysis. First, note that the coefficient of $LogFails$ in the first regression is statistically significant ($p < 0.001$). The coefficient of 1.88 means that a 1% increase in reported failure incurs an additional 1.88% increase in visits.[5] Additionally, note that the coefficient of $Full$ is also significant ($p < 0.001$). The coefficient 0.38 indicates that for $Full$ operators the number of visits per reported failure increase by 46%[6] (compared to an operator on *Basic*).

Table 5 Effect of MSP on Service Performance and Operational Costs

DV:	$Visits$	$Labor$	$Spares$
Model:	GLM with log link and Poisson distribution		
Warranty	$0.42(0.06)^{***}$	$0.06(0.09)^{ns}$	$0.13(0.25)^{ns}$
Full	$0.38(0.06)^{***}$	$0.10(0.09)^{ns}$	$0.24(0.25)^{ns}$
Partial	$0.25(0.11)^{*}$	$0.19(0.12)^{ns}$	$-0.25(0.54)^{ns}$
Basic		BASELINE	
Age	$-0.00(0.00)^{***}$	$-0.00(0.00)^{ns}$	$-0.01(0.00)^{***}$
LogFails	$1.88(0.02)^{***}$	-	-
LogVisits	-	$1.95(0.04)^{***}$	$1.77(0.06)^{***}$
Fixed Effects	Yes	Yes	Yes
Observations	23,115	23,115	22,473
Log-likelihood	-20,360	-31,211	-45,015,120

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, ^{ns} not significant

Standard Errors are clustered by equipment

If we focus on the second and third regression, we see that $LogVisits$ is highly significant ($p < 0.001$) to both *Labor* and *Spares*. The interpretation of the coefficient 1.95 for *Labor* means that a 1% increase in onsite visits incurs an expected 1.95% increase in *Labor* (similarly, the coefficient 1.77 for *Spares* corresponds to the expected percentage increase in *Spares* to a 1% increase in $Visits$). Interestingly, for both these regressions the coefficients for $Full$ is not statistically significant ($p = 0.29$ and $p = 0.34$ respectively). The interpretation is that the number of visits can potentially explain all the increases in *Labor* and *Spares*. Thus, given that the service engineer visits onsite, the manner of his repair is statistically similar across visits.

[5] The relationship between $Visits$ and $Fails$ (ignoring the other parameters) is $Log[Visits] = \beta_l Log[Fails]$. Differentiating, we get $\beta_l = \frac{dVisits}{Visits} / \frac{dFails}{Fails}$. The interpretation of this coefficient–the percentage change in $Visits$ with respect to a percentage change in $Fails$–is the elasticity of visits with respect to reported failures.

[6] The interpretation is that increases in failure rates can only partially explain the increase in onsite visits. The remaining additional increase of 46% is attributable to other factors, such as service-provider having a higher likelihood of visiting operators on $Full$ compared to *Basic*.

Figure 4 provides a diagrammatic view of the effects of MSP on various service outcomes.

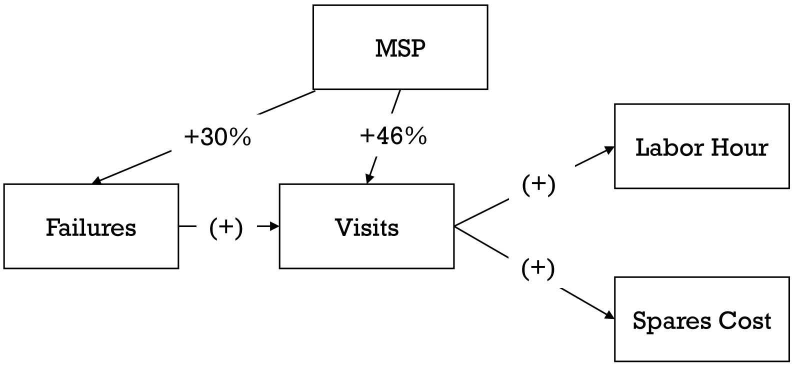

Figure 4 **A diagrammatic representation of the effect of MSP on service outcomes**

Essentially, following the arrows of the diagram, we see the MSP has two pathways of effects. First, the full-protection MSP drives a 30% higher reporting of failures by the operator (which in turn leads to more onsite visits). Second, the full-protection MSP also drives the service provider to make 46% more onsite visits given a reported failure. Altogether, these two effects lead to the large increase in onsite visits we observed in the main analysis. Finally, the increase in onsite visits directly drives increases in labor and spares cost–in fact, it fully mediates the effects we see on labor and spares. Hence, cutting down on site visits, i.e., treating problems remotely more often, will reduce service costs. In an analysis of a few months of failure data in which we obtained the description of the failure and the effect of the problem on equipment functionality (not reported here), we found that for problems of exactly the same nature, the service provider visits more often the operator on a *Full* contract rather than solving the problems remotely, as compared to an operator on a *Basic* contract. Increasing the proportion of problems resolved via remote resolution thus appears to be a feasible strategy to reduce cost (though on the assumption that it does not affect the service provider's segmentation strategy).

7. Robustness and Additional Results

7.1. Equipment Use

A fixed effects analysis depends on the assumption that the operational characteristics of each operator are stable over time. Using scan data from a small sample of 471 observed months (covering 22 pieces of equipment over time), we investigate if there are changes in operator usage patterns over time[7]. The use data–measured as number of scans in

[7] Note that the use data is also masked by a multiplication factor.

the month–were made available via a trial program where the status of the equipment is remotely tracked in real time. We plot the use data over time for a sample of six pieces of equipment in Figure 5.

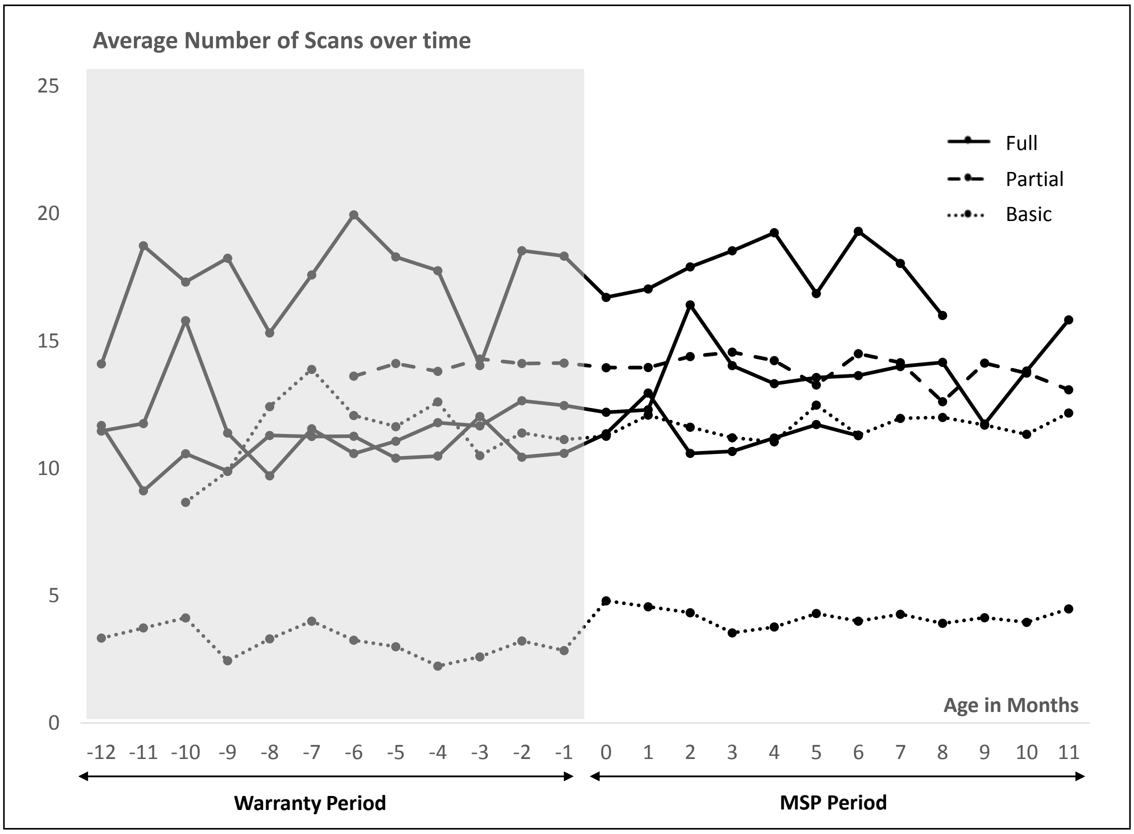

Figure 5 A plot of number of scans over time

Notice that the number of scans differs a lot from operator to operator. Yet, for the same operator, the number of scans over time seems to be very stable. Crucially, note that a change in contract type (before and after warranty ends) does not result in visually detectable changes in number of scans. We test this formally (again using a Poisson regression with bootstrapped cluster-robust errors, in the same setup as Equation 4). The results are reported in Table 6.

Note that the coefficient on *Age* is not statistically significant, meaning that there is no overall trend over time in terms of operator usage pattern, which is consistent with Figure 5 that equipment use tends to be stable over the horizon of our study. The coefficients for $Warranty$, $Full$ and $Partial$ are also not statistically significant.

Table 6 Effect of Plan on Use (GLM with log link and Poisson distribution)

DV:	*Scans*
Warranty	-0.10(0.07)^{ns}
Full (ΔC_B^F)	0.02(0.08)^{ns}
Partial (ΔC_B^P)	-0.05(0.06)^{ns}
Basic	BASELINE
Age	0.00(0.00)^{ns}
Fixed Effects	Yes
Observations	471
Log-likelihood	-828

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, ^{ns} not significant

Standard Errors are clustered by equipment

Hence, the results suggest that while equipment use may vary across operators, it tends to be fairly stable over time for each operator, providing some validity to the fixed effect assumption in the main analysis. Additionally, operators exhibit little evidence of manipulating patient scans to adjust to the MSP they are on.

7.2. Intertemporal Gaming

An additional way to test the assumptions of fixed effects (using the full dataset) is to see if the common trends of failure appear to exhibit breaks during the warranty period. Such breaks can be a result of intertemporal gaming by certain groups of operators. For example, operators who know their standard warranty will end soon and are intending to purchase only the basic plan may theoretically shift demand ahead of time just before the warranty ends.

We test for this possibility by searching for statistical evidence of systematic differences in failure patterns across different groups of operators during the warranty period. Methodologically, Angrist and Pischke (2008) propose to focus on the period prior to the treatment (i.e. before standard warranty ends), and repeat the same empirical setup pretending that the warranty period actually ended earlier than it actually did. Absent intertemporal gaming issues, the effects seen in the main model should not be seen here.

Operationally, we break the standard warranty period into two sub-periods: the first sub-period covers the beginning of the standard warranty period up until six months prior to the end of the warranty, and the second sub-period covers the remaining six months of the standard warranty. Table 7 shows the results. Note that none of the coefficients for *Full* and *Partial* showed a statistically significant value (at the 5% level) for any of the models, suggesting that operators across different MSPs are not behaving systematically differently during the warranty period in a way that would influence our estimation. The results provide further support to conclusions drawn in the paper.

Table 7 Repeating the main model assuming that the standard warranty period ended six-months earlier

DV:	*Fails*	*Visits*	*Labor*	*Spares*
Model:	GLM with log link and Poisson distribution			
Warranty	$0.15(0.03)^{***}$	$0.43(0.11)^{***}$	$0.43(0.15)^{**}$	$0.13(0.26)^{\mathrm{ns}}$
Full (ΔC_B^F)	$0.05(0.03)^{\mathrm{ns}}$	$0.17(0.11)^{\mathrm{ns}}$	$0.17(0.16)^{\mathrm{ns}}$	$0.43(0.28)^{\mathrm{ns}}$
Partial (ΔC_B^P)	$-0.02(0.07)^{\mathrm{ns}}$	$0.06(0.18)^{\mathrm{ns}}$	$0.05(0.26)^{\mathrm{ns}}$	$-0.45(0.34)^{\mathrm{ns}}$
Basic	BASELINE			
Age	$-0.01(0.00)^{*}$	$-0.01(0.00)^{***}$	$-0.02(0.01)^{***}$	$-0.01(0.00)^{\mathrm{ns}}$
Fixed Effects	Yes	Yes	Yes	Yes
Observations	10,115	10,019	10,019	9,163
Log-likelihood	-14,343	-14,447	-42,568	-25,718,120

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, $^{\mathrm{ns}}$ not significant

Standard Errors are clustered by equipment

7.3. Effect persistence

Here, we test if the effects persist over time, i.e., do the effects last after six-months, when the typical MSP has run half its course? If the effects disappear, then it may suggest that behaviors across MSP groups converge over time. Alternatively, if the effects magnify, then it may suggest the presence of positive feedback loops that accentuate a divergence in behavior across different operator groups over time.

Operationally, we test for effect persistence by re-running the main analysis with observations in the six month period just *after* warranty ends dropped. So, we focus on estimating if the effects remain after six months, when a typical MSP has run half its annual duration. Table 8 reports the results of this estimation.

Notice first that the coefficients for $Full(\Delta C_B^F)$ are all statistically significant and positive, indicating that the incentive effects of the MSP do not disappear after six months. Also, note that the coefficients for $Partial(\Delta C_B^P)$ are statistically significant and positive

Table 8 Effect of MSP on Service Costs (with six months of data immediately after warranty dropped)

DV:	*Fails*	*Visits*	*Labor*	*Spares*
Model:	GLM with log link and Poisson distribution			
Warranty	$0.34(0.06)^{***}$	$0.67(0.11)^{***}$	$0.65(0.15)^{***}$	$0.31(0.28)^{ns}$
Full (ΔC_B^F)	$0.25(0.06)^{***}$	$0.64(0.11)^{***}$	$0.68(0.15)^{***}$	$0.60(0.28)^{*}$
Partial (ΔC_B^P)	$0.22(0.09)^{**}$	$0.48(0.18)^{**}$	$0.60(0.20)^{**}$	$-0.33(0.55)^{ns}$
Basic	BASELINE			
Age	$-0.01(0.00)^{***}$	$-0.01(0.00)^{***}$	$-0.02(0.00)^{***}$	$-0.01(0.00)^{ns}$
Fixed Effects	Yes	Yes	Yes	Yes
Observations	19,204	19,158	19,158	18,464
Log-likelihood	-27,434	-26,571	-77,024	-58,653,122

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, ^{ns} not significant

Standard Errors are clustered by equipment

only for $Fails, Visits$, and *Labor*, but not for *Spares* (which is not covered by the partial-protection plan), again consistent with our main analysis. Table 9 provides a side-by-side comparison of ΔC_B^F, the effect of moving an operator from *Basic* to *Full* in the main model in Table 4 as opposed to the results in Table 8. Comparing the top and bottom rows, we can see that the mean effect sizes do not seem to diminish over time. The results suggest that behavioral changes over time, if any, occur at a fairly slow rate.

Table 9 Comparison of estimates for ΔC_B^F, the effect of moving an operator from *Basic* to *Full*

DV:	*Fails*	*Visits*	*Labor*	*Spares*
Main Model:	30%	80%	69%	102%
After Six Months:	28%	90%	97%	82%

8. Discussion and Conclusion

The present paper establishes that, when properly controlling for contract selection, a fixed-fee full-protection plan leads to more failures and higher service costs compared to a time-and-materials plan in the context of medical equipment maintenance. Further, our mediation analysis shows that both the operator and the service provider contribute to the increase in service costs. Our results thus indicate that MSP structure impacts the behavior of both the operator and the service provider.

Using a smaller sample of usage data, we find no evidence of operator moral hazard arising from changes in equipment usage, which is consistent with empirical findings by Ning et al. 2014 in the printing industry. One possible reason for this is that the operator's utility is directly derived from usage, so that the operator has the incentive to utilize equipment as high as demand allows, irrespective of the contract structure, and thus lessening the scope at which usage is manipulated. This also suggests that quality of operational care, which is hard to observe and quantify, is the key information asymmetry affecting the performance of the maintenance service chain.

As for the service provider, we consistently see (across the three types of contracts we have) that he expends more resources when he bears the costs for those resources. Hence, the service provider's moral hazard appears to be driven less by cost concerns but rather by concerns of providing differentiated levels of onsite responsiveness to operators on different plans.

Our work thus provides the first empirical evidence that a basic time-and-materials plan can outperform a full-protection plan in after-sales maintenance, implying that service providers should be wary about assuming more responsibility over equipment failure outcomes. While Roels et al. 2010 have pointed to this theoretical possibility, currently no empirical evidence exists to support it. In fact, Guajardo et al. 2012, in an empirical study of aircraft engine maintenance, showed evidence to the contrary, i.e., having the service provider assume responsibilities for both equipment failure and repair costs leads to less failure events.

Taken together, our results establish that the full-protection plan results in an unambiguous decrease in maintenance service chain value relative to the basic (time- and-material) plan. But while we can draw conclusions on the service chain value, the lack of revenue data, which is highly commercially sensitive, prevented us from delineating the implications for the service provider and the operators' profits. An interesting avenue of future research is to analyze how such value is distributed across the service provider and the operator. Given the rising expenditures related to diagnostic imaging in the healthcare industry, it is important that we better understand the sources of inefficiencies and whether alternative contract structures exist that can Pareto improve existing ones.

Finally, future research could also examine the dynamics of contract selection and moral hazard over the equipment lifecycle. While more than 90% of operators in our setting renew with the same type of plans, there is evidence that plan switching becomes more prevalent toward the end of the equipment's life, indicating changes to the nature of the interactions as time progresses. Examining such dynamics can uncover insights into how firms can create comprehensive, or even dynamically optimized, contract and service management spanning over the equipment life.

Acknowledgments

We are grateful for the support of Siemens Healthcare Customer Services. In particular, we thank Alfred Fahringer, Marc Muehlen, and Rajneesh Moudgil from Siemens Regional Headquarters Asia-Australia for insightful discussions and input.

References

Abbring, J. H., J. J. Heckman, P.-A. Chiappori, J. Pinquet. 2003. Adverse Selection and Moral Hazard In Insurance: Can Dynamic Data Help to Distinguish? *Journal of the European Economic Association* **1**(2-3) 512–521.

Angrist, J. D., J. S. Pischke. 2008. *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.

Banerjee, A. V., E. Duflo. 2000. Reputation effects and the limits of contracting: A study of the indian software industry. *Quarterly Journal of Economics* 989–1017.

Baron, R. M., D. A. Kenny. 1986. The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of personality and social psychology* **51**(6) 1173.

Cachon, G. P. 2003. Supply chain coordination with contracts. *Handbooks in Operations Research and Management Science* **11** 227–339.

Cameron, A. C., P. K. Trivedi. 2009. *Microeconometrics using Stata*. Stata Press, College Station, TX.

Chiappori, P.-A., B. Salanié. 2003. Testing contract theory: A survey of some recent work. *Econometric Society Monographs* **35**.

Corts, K. S., J. Singh. 2004. The Effect of Repeated Interaction on Contract Choice: Evidence from Offshore Drilling. *Journal of Law, Economics, and Organization* **20**(1) 230–260.

Cui, R., G. Allon, A. Bassamboo, J. A. Van Mieghem. 2013. Information Sharing in Supply Chains: An Empirical and Theoretical Valuation. Manuscript submitted for publication.

Deloitte Research. 2006. The Service Revolution in Global Manufacturing Industries.

ECRI Institute. 2013. Healthcare Product Comparison System.

Fuchs, V. R., I. H. Leveson. 1968. *The Service Economy*. National Bureau of Economic Research.

Gopal, A., K. Sivaramakrishnan, M. S. Krishnan. 2003. Contracts in offshore software development: An empirical analysis. *Management Science* **49**(12) 1671–1683.

Guajardo, J. A., M. A. Cohen, S.-H. Kim, S. Netessine. 2012. Impact of Performance-Based Contracting on Product Reliability: An Empirical Analysis. *Management Science* **58**(5) 961–979.

Jain, N., S. Hasija, D. G. Popescu. 2013. Optimal contracts for outsourcing of repair and restoration services. *Operations Research* **61**(6) 1295–1311.

Kalnins, A., K. J. Mayer. 2004. Relationships and hybrid contracts: An analysis of contract choice in information technology. *Journal of Law, Economics, and Organization* **20**(1) 207–229.

Kawaja, J., D. Durmis. 2005. MRI coil failures: Decreasing occurrence and cost. *Thermo USCS Company Knowledge Base* .

Kim, S.-H., M. A. Cohen, S. Netessine. 2007. Performance Contracting in After-Sales Service Supply Chains. *Management Science* **53**(12) 1843–1858.

Nagarajan, M., G. Sošić. 2008. Game-theoretic analysis of cooperation among supply chain agents: review and extensions. *European Journal of Operational Research* **187**(3) 719–745.

Ning, J., V. Babich, J. Handley, J. Keppo. 2014. Empirical Study on Managed Print Services Pricing. Manuscript submitted for publication.

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

Novak, S., S. Stern. 2009. Complementarity among vertical integration decisions: Evidence from automobile product development. *Management Science* **55**(2) 311–332.

Oliva, R., R. Kallenberg. 2003. Managing the transition from products to services. *International Journal of Service Industry Management* **14**(2) 160–172.

Plambeck, E. L., S. A. Zenios. 2000. Performance-based incentives in a dynamic principal-agent model. *Manufacturing & Service Operations Management* **2**(3) 240–263.

Ramdas, K., T. Randall. 2008. Does Component Sharing Help or Hurt Reliability? An Empirical Study in the Automotive Industry. *Management Science* **54**(5) 922–938.

Roels, G., U. S. Karmarkar, S. Carr. 2010. Contracting for Collaborative Services. *Management Science* **56**(5) 849–863.

Sawhney, M., S. Balasubramanian, V. Krishnan. 2004. Creating growth with services. *MIT Sloan Management Review* **45**(2) 1–17.

Susarla, A. 2012. Contractual Flexibility, Rent Seeking, and Renegotiation Design: An Empirical Analysis of Information Technology Outsourcing Contracts. *Management Science* **58**(7) 1388–1407.

Susarla, A., R. Subramanyam, P. Karhade. 2010. Contractual Provisions to Mitigate Holdup: Evidence from Information Technology Outsourcing. *Information Systems Research* **21**(1) 37–55.

Terwiesch, C., Z. J. Ren, T. H. Ho, M. A. Cohen. 2005. An Empirical Analysis of Forecast Sharing in the Semiconductor Equipment Supply Chain. *Management Science* **51**(2) 208–220.

Wooldridge, J. M. 2010. *Econometric analysis of cross section and panel data*. MIT Press, Cambridge, MA.

Zhou, Y.-P., Z. J. Ren. 2011. Service outsourcing. *Wiley Encyclopedia of Operations Research and Management Science* .

Recent ESMT Working Paper

	ESMT No.
Designing luxury experience Vadim Grigorian, Pernod Ricard Francine Espinoza Petersen, ESMT	14-04
Financing capacity investment under demand uncertainty Francis de Véricourt, ESMT Denis Gromb, INSEAD	14-03
Government guarantees and bank risk taking incentives Markus Fischer, Goethe-Universität Frankfurt am Main Christa Hainz, ifo Institute for Economic Research Jörg Rocholl, ESMT Sascha Steffen, ESMT	14-02
What do patent-based measures tell us about product commercialization? Evidence from the pharmaceutical industry Stefan Wagner, ESMT Simon Wakeman, ESMT	14-01
Technology commercialization strategy in a dynamic context: Developing specialized complementary assets in entrepreneurial firms David H. Hsu, The Wharton School, University of Pennsylvania Simon Wakeman, ESMT	11-02 (R4)
The equivalence of bundling and advance sales Alexei Alexandrov, Consumer Financial Protection Bureau, Washington DC Özlem Bedre-Defolie, ESMT	13-11
Hidden efficiencies: On the relevance of business justifications in abuse of dominance cases Hans W. Friederiszick, E.CA Economics, ESMT Linda Gratz, E.CA Economics	13-10
Overcoming localization of knowledge: The role of professional service firms Stefan Wagner, ESMT Karin Hoisl, University of Munich Grid Thoma, University of Camerino	13-09

ESMT
European School of Management and Technology
Faculty Publications
Schlossplatz 1
10178 Berlin
Germany

Phone: +49 30 21231-1279
publications@esmt.org
http://www.esmt.org/faculty-research/research-and-development/publications