



HAL
open science

Deterrence and learning effects in insurance fraud audits

Reda Aboutajdine

► **To cite this version:**

Reda Aboutajdine. Deterrence and learning effects in insurance fraud audits. Economics and Finance. Institut Polytechnique de Paris, 2019. English. NNT : 2019IPPAX016 . tel-03160538

HAL Id: tel-03160538

<https://theses.hal.science/tel-03160538v1>

Submitted on 5 Mar 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 2019IPPAX016

Thèse de doctorat



Deterrence and learning effects in insurance fraud audits

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École Polytechnique

École doctorale n°626 Ecole Doctorale de l'Institut Polytechnique
de Paris (ED IP Paris)
Spécialité de doctorat : Sciences Économiques

Thèse présentée et soutenue à Palaiseau, le 20 Décembre 2019, par

REDA ABOUTAJDINE

Composition du Jury :

Olivier Gossner Directeur de recherche, CNRS	Président
Eduardo Perez-Richet Professeur Associé (HDR), Sciences Po	Rapporteur
Jean Pinquet Maître de Conférences (HDR), Université Paris-Nanterre	Rapporteur
Wanda Mimra Professeur Associée, IÉSEG	Examineur
Pierre Picard Professeur, membre du Haut Collège, École Polytechnique	Directeur de thèse

Remerciements

Tout d'abord, j'aimerais exprimer ma sincère et profonde reconnaissance envers Pierre Picard, mon directeur de thèse, pour son encadrement rigoureux et son soutien indéfectible pendant cette thèse. Coécrire avec lui aura été une expérience fort enrichissante, à la fois intellectuellement et humainement. J'ai été particulièrement inspiré par son enthousiasme sans limite pour la recherche et par l'étendue et la profondeur de sa connaissance des sciences économiques. J'espère que nous continuerons à travailler ensemble.

Mes remerciements vont également à Olivier Gossner, Wanda Mimra, Eduardo Perez-Richet et Jean Pinquet pour m'avoir fait l'honneur de participer à mon jury de thèse. J'espère que la lecture de ce manuscrit aura réussi à les intéresser à l'analyse économique de la fraude à l'assurance.

J'ai eu la chance durant ma thèse de bénéficier d'environnements de recherche exceptionnels, d'abord au Laboratoire d'Économétrie de l'École Polytechnique, puis au CREST. Je remercie particulièrement Pierre Boyer pour sa disponibilité et l'intérêt qu'il porte à tous les doctorants sans distinction et Isabelle Méjean pour sa bienveillance. Merci aussi à François Pannequin et Thibault Richard pour les nombreuses discussions sur l'assurance comportementale. J'ai aussi pris énormément de plaisir à donner des cours et TD dans le cadre du Master in Economics, face à des élèves exceptionnels, et dont certains ont déjà rejoint les rangs du CREST en tant que doctorants (ô vieille ennemie !). J'ai notamment beaucoup apprécié de travailler avec Xavier d'Haultfoeuille et Thierry Kamionka pour les TD d'Économétrie 1. Je tiens à remercier Lyza Racon, Eliane Madeleine et Weronika Leduc pour leur gentillesse et leur patience face à mes longs et ennuyeux mails administratifs. Quant à Sri Srikandan, le premier chapitre de cette thèse n'aurait jamais pu exister sans son aide, et pour cela je lui suis reconnaissant. Enfin, je remercie les membres du CASD pour le travail remarquable et fondamental qu'ils font au quotidien.

J'ai également bénéficié d'un environnement professionnel stimulant au sein des différentes équipes d'IBM France. Travailler au sein de l'équipe Solon m'aura permis de confronter mes

idées théoriques et abstraites de doctorant aux réalités du terrain et, en retour, de m'en inspirer pour avoir des idées de sujet de recherche. J'ai une pensée particulière pour Emmanuelle Dulac, dont la capacité à comprendre intuitivement et rapidement mes explications trop techniques de méthodes économétriques à tester m'impressionnera toujours. Mon expérience au sein du France Lab aura été très enrichissante, dans un environnement plus technique. Je tiens en particulier à exprimer ma reconnaissance envers Christian de Sainte Marie pour son accompagnement pendant cette thèse et son dévouement aux doctorants du France Lab. Je garde un excellent souvenir des discussions de bureau avec Thomas Baudel, donc la culture générale ne cessera de m'impressionner. J'ai aussi une pensée particulière pour les doctorants du France Lab, Hamza, Penelope, Karim, Oumaima et Olivier. Enfin, dans le cadre du projet Solon, j'ai eu l'occasion de collaborer avec des équipes de PRO BTP. Je remercie particulièrement Frédéric Estroumza et Pierre-Henri Darnault qui m'ont notamment accueilli plusieurs fois à Cagnes-sur-Mer pour m'aider à élucider les nombreuses questions que j'avais sur la fraude à l'assurance santé.

Un doctorat, c'est aussi une famille de doctorants. Je pense tout d'abord à mon senpai from the same supervisor, Alexis Louaas, mais aussi à mes acolytes du bureau Polytechnique, Anais Goburdhun, Margot Hovsépian, Hugo Molina, Anasuya Raj et Rafael Veiel, à ceux du bureau CREST, Alexis Larousse, Louis Pape et Antoine Valtat, et à celui de la coloc, Jérôme Trinh. Enfin, un grand merci pour les foots, les soirées raclettes et les KaraoCrest à Rémi Avignon, Antoine Bertheau, Aurélien Bigo, Guidogiorgio Bodrato, Pauline Carry, Arthur Cazaubiel, Gwen-Jiro Clochard, Pierre-Edouard Collignon, Morgane Cure, Thomas Delemotte, Antoine Ferey, Christophe Gaillac, Germain Gauthier, Lucas Girard, Morgane Guignard, Aymeric Guidoux, Etienne Guigue, Yannick Guyonvarch, Jeremy Hervelin, Tomas Jagelka, Dimitrios Karoukis, Jeremy L'Hour, Sébastien Laffitte, Thomas Larrieu, Clémence Lenoir, Alicia Marguerie, Esther Mbih, Alfonso Montes, Martin Mugnier, Elio Nimier, Ivan Ouss, Bérengère Patault, Fabien Perez, Emilie Sartre, Clémence Tricaud, Ao Wang, Julien Combe, Jeanne Commault, Esther Delbourg, Sebastian Franco Bedoya, Arnaud Goussebaile, Alena Kotelnikova, Guilhem Lecouteux, Gwenaël Roudaut.

Je n'oublie évidemment pas mes secondes familles, les amis d'antan et d'aujourd'hui, Collégiens, Cartésiens, Jansonniens, Télécommiens, PSEiens ou autres. J'en ai dressé une liste merveilleuse et exhaustive, mais "la marge est trop petite pour la contenir".

J'ai une pensée très affectueuse pour ma première famille, ma mère, Atiqa Rifai, et mes soeurs, Fatine et Samya. La période auxquelles correspond cette thèse aura coïncidé avec certains moments

particulièrement douloureux pour notre famille, mais nous avons su les dépasser ensemble. Ce dont je n'ai aucun doute, c'est que sans l'amour et l'affection inconditionnels que je vous porte, et dont j'ai bénéficié en retour de votre part, toute cette entreprise n'aurait jamais été possible, et je n'aurais jamais pu mener cette thèse à son terme.

Enfin, c'est avec beaucoup d'appréhension que j'en arrive à écrire ces derniers mots pour clore ces remerciements. Papa, cette thèse t'est dédiée. Je ne sais pas si je parviendrai un jour à me remettre de ta disparition si soudaine et si brusque, au coeur même d'une thèse qui t'est indissociable. Alors même que mon intérêt pour la recherche et ses raisonnements rigoureux émanait indubitablement de ta propre passion et des discussions interminables que nous avons eu à ce sujet tout au long de notre vie commune, voilà que tu n'es plus là pour être témoin de cet aboutissement. La peine que suscite en moi ton absence en ce moment est ineffable, et elle restera toujours une déchirure pour moi. Je me résous juste à espérer que le fait de m'être entêté à finir cette thèse contre vents et marées permettra d'atténuer un tant soit peu cette douleur. Car finalement, faire preuve d'entêtement et de résilience demeure le plus bel hommage que je pourrais te rendre.

Résumé

La fraude est un problème majeur des marchés d'assurance : elle conduit à des accroissements de coûts et donc à une hausse des primes. De ce fait, elle réduit l'efficacité des mécanismes assurantiels de partage des risques. Le 1er Octobre 2019, la Caisse nationale de l'assurance maladie a publié son « Bilan 2018 des actions contre la fraude et actions de contrôles ». Selon ce rapport, le montant des fraudes détectées et sanctionnées entre 2014 et 2018 avoisine 1,2 milliard d'euros. En outre, la fraude à l'assurance est un phénomène endémique qui ne se cantonne pas au domaine de la santé. Elle touche des secteurs aussi divers que l'assurance automobile, l'assurance incendie ou la prévoyance.

L'audit est indubitablement le principal outil de lutte anti-fraude à l'assurance. En vérifiant la validité des demandes de remboursement qui lui sont soumises, l'assureur peut recouvrer les indus. Mais le rôle de l'audit ne se résume pas à ce recouvrement. Tout d'abord, il fait office de menace pour les fraudeurs et possède un effet dissuasif à même de décourager, partiellement ou totalement, leurs activités malhonnêtes. Il a aussi un rôle informationnel puisqu'il permet de mieux comprendre les mécanismes sous-jacents et de mieux cibler les audits futurs. Or l'audit constitue un coût sûr pour des bénéfices qui eux ne sont pas garantis. Il est donc primordial d'analyser rigoureusement tous les avantages liés à ses différents rôles afin d'élaborer des stratégies d'audit optimales et d'en évaluer l'efficacité.

Partant de ce constat, cette thèse en étudie les fonctions dissuasive et informationnelle. Elle s'intéresse notamment aux effets dissuasifs dans le cadre du développement récent de méthodes d'apprentissage statistique et de la démocratisation des données massives (Big Data). Ce développement s'est traduit par le déploiement de solutions technologiques de lutte anti-fraude pour un audit plus précis. On s'intéresse également aux mécanismes de collusion par lesquels des assurés et des prestataires (opticiens, garagistes, etc.) peuvent s'entendre pour frauder. Dans la mesure où ces prestataires traitent souvent plusieurs assurés d'un même assureur, les interactions prestataire-assureur sont fréquentes et répétées, ce qui renforce le rôle informationnel de l'audit.

Le premier chapitre de cette thèse est une étude empirique des effets dissuasifs de l'audit, fruit d'une collaboration avec IBM France et PRO BTP, dans le cadre du déploiement de la solution anti-fraude Solon. Cette analyse montre que les audits subis par un opticien ont pour effet de réduire sa fraude future. Plus spécifiquement, l'effet dissuasif est d'autant plus fort que la lettre d'audit est crédible, soulignant l'importance de l'engagement (« *commitment* ») dans la lutte contre la fraude.

Les deuxième et troisième chapitres étudient un problème d'audit dynamique où l'auditeur interagit de façon répétée avec des prestataires et peut apprendre des résultats de son audit. Le deuxième chapitre utilise un modèle à deux périodes pour mettre en évidence cet effet d'apprentissage dont la conséquence est qu'il est optimal d'auditer plus intensément au début de la relation. Le troisième chapitre étend ce modèle à un nombre arbitraire ou infini de périodes, et montre que l'audit optimal est d'autant plus intense qu'il reste un nombre important d'interactions. L'intuition réside dans le fait que le coût supplémentaire d'un audit présent a des répercussions positives sur toutes les périodes suivantes.

Enfin, le quatrième chapitre réunit les mécanismes de dissuasion et d'apprentissage dans un même modèle dynamique. Il montre que l'apprentissage transforme la dissuasion en une menace intertemporelle. En d'autres termes, un audité sera dissuadé aujourd'hui car il risque de détériorer sa réputation future s'il se fait attraper en train de frauder.

Mots clés: Fraude, Assurance, Dissuasion, Apprentissage, Interactions répétées

Contents

Remerciements	i
Résumé	v
Introduction	1
1 Every Claim You Make I'll be Watching You! The Deterrence Effects of Insurance Claim Audits	5
1.1 Introduction	7
1.2 Research Setting and Empirical Strategy	11
1.3 Empirical Results	22
1.4 Discussion: Limits and Extensions	26
1.5 Conclusion	30
2 Preliminary Investigations for Better Monitoring: Learning in Repeated Insurance Audits	31
2.1 Introduction	33
2.2 Single Period Auditing	35
2.3 Two-Period Auditing: The Learning Effect	37
2.4 Variable Claim Value	46
2.5 Conclusion	50
2.6 Appendix	52
3 Should I Stalk or Should I Go? An Auditing Exploration/Exploitation Dilemma	61
3.1 Introduction	63
3.2 The Model	66
3.3 Exploration vs Exploitation: Auditing to Separate the Wheat from the Chaff . .	74

3.4	The Extents of Learning	83
3.5	Conclusion	87
3.6	Appendix	89
4	You'll Give Yourself a Bad Name: Reputation Effects in Repeated Audits	101
4.1	Introduction	103
4.2	Model	105
4.3	Reputation-based Deterrence	113
4.4	Auditing with Deterrence and Learning Effects and the Restless Bandit Problem	123
4.5	Conclusion	124
4.6	Appendix	126
	Conclusion	141
	<i>Bibliographie</i>	151
	<i>Liste des tableaux</i>	153
	<i>Liste des figures</i>	156

Introduction

Insurance fraud represents a severe and pervasive problem for all types of non-life insurance.¹²³⁴⁵ Not even disaster insurance is spared by scams.⁶ According to the Coalition Against Insurance Fraud (CAIF), fraud corresponds to an annual loss of \$80 billion across all lines of insurance,⁷ and common estimates put fraud at approximately 10% of incurred losses.⁸ By generating additional costs, fraud increases the premia, as insurers pass on costs to policyholders.⁹ The FBI estimates the increase in premia for non-health insurance to be between \$400 to \$700 for the average US family.¹⁰ Therefore, fraud reduces the efficiency of risk sharing mechanisms. In some extreme cases, it may even put insurance markets at the risk of collapsing. With ageing populations, healthcare insurance and workers compensation will see the corresponding premia increase, which, combined with the increase due to fraud, might drastically reduce demand for insurance. The same situation might happen for catastrophic insurance because of climate change. In addition, the fact that insurers pass on costs to the insured can be attributed to a form of laxness from their part.¹¹ As a consequence, it may deteriorate the already bad reputation insurers have, and encourage more fraud.¹²

Yet, while this might be true, there are other plausible explanations for why fraud is not systematically addressed by insurers. Indeed, fraud is a difficult issue to tackle and legal

¹ Auto insurance <http://www.insurancefraud.org/downloads/InsuranceResearchCouncil02-15.pdf>

² Healthcare https://www.ameli.fr/fileadmin/user_upload/documents/Controles_et_Fraudes.pdf

³ Disability <https://www.bustathief.com/disability-fraud-disability-scam/>

⁴ Workers Compensation <https://www.insurancefraud.org/scam-alerts-workers-compensation.htm>

⁵ Property https://www.insurancefraud.org/downloads/IFPA_Homeowners_Insurance_Fraud.pdf

⁶ <https://www.fbi.gov/stats-services/publications/insurance-fraud>

⁷ <http://www.insurancefraud.org/statistics.htm>

⁸ e.g., [Derrig \(2002\)](#) and <https://www.iii.org/article/background-on-insurance-fraud>

⁹ <https://www.fraudny.com/youpay>

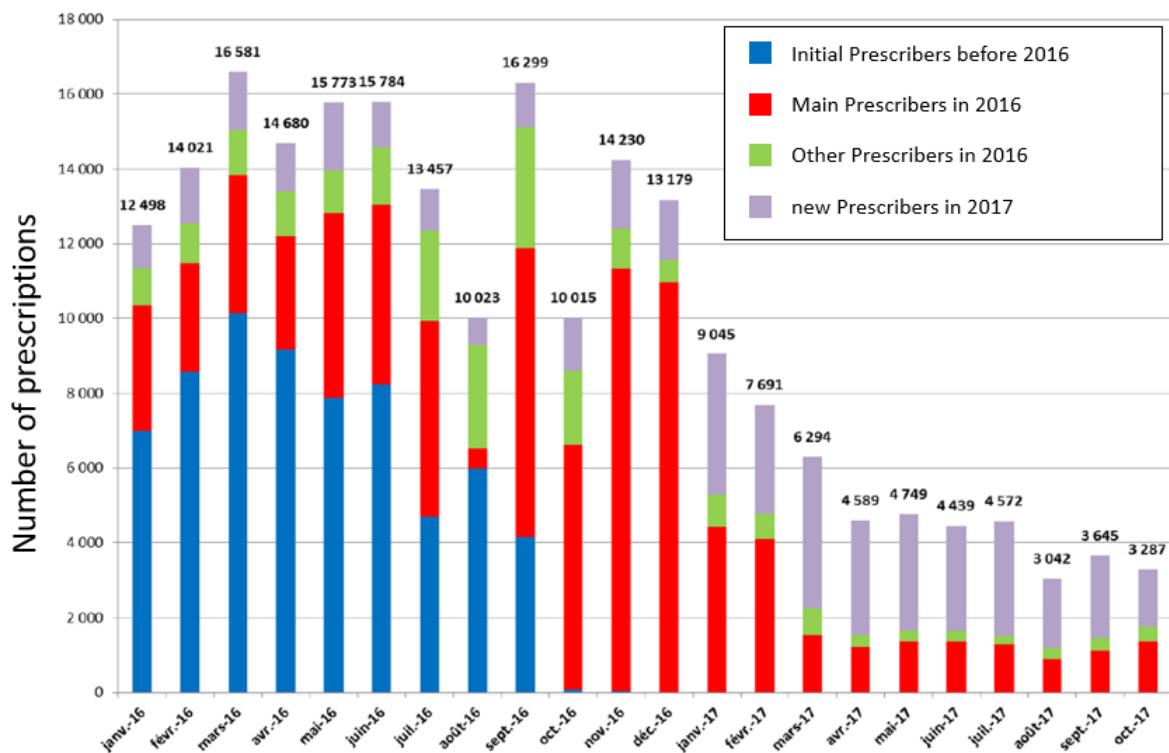
¹⁰ <https://www.fbi.gov/stats-services/publications/insurance-fraud>

¹¹ <https://www.vox.com/science-and-health/2019/7/19/20698355/health-insurance-scam-united-cigna-aetna-southwest>

¹² <https://newsroom.accenture.com/industries/insurance/one-fourth-americans-say-its-acceptable-to-defraud-insurance-companies-accenture-survey-finds.htm>

prosecution costs can end up being larger than the amount at stake.¹³¹⁴ As a consequence, insurers exclude claims under a threshold amount from prosecution. Furthermore, defrauders can build sophisticated fraud schemes and adapt to the insurer’s counterfraud efforts. For example,¹³ David Williams, a Texan kinesiologist, acquired a medical license he was not entitled to get, and billed insurers in increments of \$300. These small claims did not make the prosecution threshold and ended adding up to about \$25 million over more than four years. Figure 1 shows how an organized medicine drug trafficking network adapted to the French CNAM’s¹⁵ efforts by changing prescribing doctors.

Figure 1 – ADAPTATION OF DEFRAUDERS AND CHANGE IN DRUG PRESCRIBERS (SOURCE CNAM²)



In order to address the costliness of prosecution and tackle the insurance fraud problem appropriately, it is therefore necessary to either increase the revenues from auditing and/or reduce its cost.

First, it is fundamental to better identify the returns from auditing. When considering the costs and benefits from auditing, insurers often reduce auditing benefits to the immediate recovered illegitimate reimbursements after prosecution, conditionally on the latter being successful. Yet,

¹³ <https://www.propublica.org/article/we-asked-prosecutors-if-health-insurance-companies-care-about-fraud-they-laughed-at-us>

¹⁴ <https://www.claimsjournal.com/news/national/2014/08/14/253331.htm>

¹⁵ Caisse Nationale d’Assurance Maladie.

auditing may involve at least two additional mechanisms that influence generated revenues. The first mechanism is deterrence, whereby sustained auditing represents a threat for potential defrauders, who consequently respond by reducing their defrauding efforts. The second mechanism is a learning one, as auditing also allows the insurer to gather some information about auditees and defrauding methods. This information allows him to refine his auditing techniques and better target inspections in the future, to only engage auditing costs when it is profitable. By accounting for both mechanisms, auditing may turn out to be profitable in situations where it is ordinarily believed not to, and therefore encourage initially reluctant insurers to address fraud rather than pass on the related costs to policyholders.

Second, it is necessary to understand how recent developments in big data analytics can help mitigate the costliness of auditing problem. The use of these methods in the context of insurance fraud is spreading quickly and proving to be an efficient tool. According to a survey by the CAIF,¹⁶ by 2016, 75% of surveyed insurers relied on anti-fraud technology for claims fraud detection. The perceived benefits can be separated into two categories: on the one hand, anti-fraud technology allows to handle more claims in the same timespan (mentioned benefits include “more referrals”, “enhanced reporting” and “improved investigator efficiency”), and, on the other hand, to have a more refined analysis of claim suspicion (“higher quality referrals”, “uncovering complex or organized fraud” and “better understanding of referrals”).

Based on this premise, this thesis studies the deterrent and learning roles of auditing. It examines the deterrent effects in the context of the recent development of machine learning techniques and the democratization of Big Data. As a consequence, technological solutions allowing for more accurate auditing have been deployed. This thesis also analyzes collusion mechanisms through which policyholders and service providers (opticians, mechanics, etc.) may collaborate to defraud. Given that these service providers usually serve several policyholders of the same insurer, provider-insurer interactions are frequent and repeated, which strengthens the learning role of auditing.

The first chapter is an empirical assessment of the deterrence effects of auditing. It was conducted in collaboration with IBM France and PRO BTP, in the context of the deployment of the Solon counter-fraud solution. This assessment shows that incurred audits decrease an optician’s subsequent fraud. More specifically, this deterrence effect is stronger the more credible the audit letter, underlining the importance of commitment in counter-fraud efforts.

The second and third chapters, co-authored with Pierre Picard, examine a dynamic auditing

¹⁶http://www.insurancefraud.org/downloads/State_of_Insurance_Fraud_Technology2016.pdf

problem where the auditor interacts repeatedly with service providers, and can learn from auditing outcomes. The second chapter relies on a two-period model to show the existence of a learning effect, which consequence is that it is optimal to audit more at the beginning of the relationship. The third chapter extends this model to an arbitrary or infinite number of periods, and shows that the further the horizon, the larger optimal auditing. Intuition stems from the fact that the cost of a current audit has a positive impact on all future periods proceeds.

Finally, the fourth chapter, co-authored with Pierre Picard, sets a model with both deterrence and learning mechanisms. It reveals a reputation based deterrence effect, whereby learning turns deterrence effects from a static threat to an intertemporal one. In other words, an auditee will be deterred in the present because of the risk of seeing his reputation deteriorate if he gets caught defrauding.

Chapter 1

Every Claim You Make I'll be Watching You! The Deterrence Effects of Insurance Claim Audits ¹

¹This work has benefited from a collaboration with teams from IBM France and PRO BTP to whom I am thankful. I am grateful to Pierre-Henri Darnault, Frédéric Estroumza, Lucas Girard, Yannick Guyonvarch, Alexis Louaas, Louis Pape, Bérengère Patault, Pierre Picard and Anasuya Raj for helpful discussions and comments. I am also grateful to Sri Srikandan for his help to make data available. This work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the “Investissements d’Avenir” program (reference: ANR-10-EQPX-17 - Centre d’accès sécurisé aux données – CASD). Funding from FX-Conseil, IBM France and Labex ECODEC is gratefully acknowledged.

Abstract

Fraudulent behavior represents a serious threat to the viability of insurance markets, and Special Investigative Units rely on auditing to tackle this problem. Rather than the recovered amounts from audited fraudulent claims, it is the deterrence effects of auditing that are believed to rationalize the non-negligible monitoring costs. Using individual-level claim submission/investigation count panel data from a health insurer, and a Poisson two-way fixed effects estimator, I provide evidence of this deterrence effect and find that each additional incurred audit decreases the auditee's current submissions by 3%. I also find that the more credible the threat, the stronger the impact of auditing. Audits that request supporting documents reduce submissions by 4.2%, while simple warning letters have no impact. This confirms the importance of commitment and credibility in counter-fraud measures.

JEL Classification Numbers: C23, G22, K42.

Keywords: Insurance fraud, Auditing, Deterrence, Unobserved effects, Count data.

1.1 Introduction

Insurance fraud poses a serious threat to insurance markets: with a share of 10%² of claims believed to be fraudulent, fraud raises the insurer's estimation of the probability of an accident and increases its costs, thus mechanically increasing the *premia*. While the recent development and adoption of machine learning techniques for fraud detection have helped increase the accuracy of auditing and lower its costs, these costs remain too high to rationalize the use of auditing solely as an undue indemnity recovery tool. Instead, Special Investigative Units (SIU) experts invoke the deterrence role of auditing: since they face the risk of being caught and incur some damage (direct financial penalty, exclusion from the insurer's network, reputational damage, ...), rational defrauders are expected to reduce their defrauding efforts. However, defrauders are never perfectly rational and many behavioral mechanisms can bias their expected reactions to auditing, sometimes even implying an increase in post-audit fraud. For example, because of the gambler's fallacy,³ an auditee could expect the short term post-audit probability of investigation to be close to zero, and try to cover the incurred damages through additional fraud. Understanding and quantifying the magnitude of audit deterrence effects is central to the improvement of counter-fraud measures.

In this paper, I empirically assess the deterrence effects of audits in the context of health insurance fraud, more specifically for vision insurance. Using industrial panel data from a French mutual insurer, I study the impact of the number of incurred audits in the past on the number of current claim submissions. More specifically, the data covers the period following the implementation of new counter-fraud measures in the context of a collaboration between the insurer, PRO BTP, and a major tech firm, IBM France. These measures correspond to the use of advanced statistical detection systems and the rationalization of audit targeting. The audits consist in different types of letters that range from a simple warning to litigation.

I use a Poisson model with two-way fixed-effects to address the identification issues of the problem. Firstly, the Poisson specification accounts for the discrete nonnegative nature of the data, where a fairly large share of the number of submitted claims per individual/period is equal to zero. The Poisson specification is also immune to the incidental parameter problem inherent to the use of individual fixed effects in short panels, and is robust to distributional misspecification under weak assumptions. Secondly, auditing is endogenous in that it targets defrauders who

²Derrig (2002)

³Mittone (2006) finds, in a tax evasion experiment, that the immediate post-audit rounds display a systematic increase in evasion. He labels it the "bomb crater effect" and explains that "the subjects decides to evade immediately after a fiscal audit because they believe that it cannot happen twice in the same place (time)".

submit more claims on average. The determinants of fraud are usually intrinsic unobservable features such as morality, technical ability (fraud is a complex process), marketing ability, etc., so individual fixed-effects allow us to correct for this selection bias. Finally, auditing is also endogenous with regards to timing as audits occur right after a suspicious activity has been spotted, and fraud is seasonal.⁴

I provide evidence that the number of undergone audits before a given period has a negative impact on the number of submitted claims at the considered period. I find that each additional audit decreases the claim count by 3% on average. I argue that two mechanisms can explain this result: first, the Service Provider's belief about the *general probability* of an audit is updated upwards after every audit, independently of the audit outcome, thus leading to a lower level of fraud intensity in the subsequent period equilibrium. Second, if the audit reveals fraud, catching a Service Provider (SP) defrauding gives information to the SIU who identifies the SP as being of the dishonest type. It is akin to a reputational damage that increases the *SP-specific probability* of being audited.

I also find that different types of audits have different levels of deterrence effects. Simple warning letters have almost no impact while audit letters requesting supporting documents induce a strong decrease (-4.2%) in subsequent claim submissions. I interpret this result in terms of commitment: simple warnings are perceived as an empty threat, while document requests are more credible because of how specific they are, and show the insurer's commitment to follow up on the threat.

This paper contributes to the large literature on criminal activity and deterrence effects. Following the seminal contributions of [Becker \(1968\)](#) and [Ehrlich \(1973\)](#), an extensive empirical literature has investigated the relationship between crime rates and law enforcement measures ([Craig \(1987\)](#), [Trumbull \(1989\)](#)). In particular, [Cornwell and Trumbull \(1994\)](#) rely on panel data to be able to control for unobserved heterogeneity, and present empirical evidence that the ability of the criminal justice to deter crime is much weaker than previous results indicate. However, they deplore the unavailability of individual level data as a limit to the assessment of the deterrence effects of law enforcement on crime. This, on the other hand, is not a problem for the study of tax evasion and insurance fraud, where both tax authorities and insurers have access to granular data about the protagonists' behavior.

⁴e.g, for automobile insurance fraud, [Picard and Wang \(2015\)](#) note that “the concentration of claims during the last months of the policy year is compatible with seasonal fluctuations in the number of claims over the calendar year, with peaks during vacation months (January, March, July and December)” and [Dionne and Wang \(2013\)](#) find that “total theft claims of no-deductible endorsement are more common during the early months of the contract year”.

The empirical literature on tax evasion and audits is remarkably rich. Following the seminal paper by [Allingham and Sandmo \(1972\)](#), several experimental and empirical contributions have focused on the behavioral responses to taxpayer audits. [Mittone \(2006\)](#) conducts an experimental study revealing a 'bomb-crater effect' whereby taxpayers cheat more following an audit. [Kleven et al. \(2011\)](#) use a field experiment in Denmark to examine the effects of two forms of threats, prior (actual) audits and threat-of-audit letters, on subsequent reporting, and show that while both have a positive effect on the subsequent declared amounts, i.e., a deterrence effect, the effect is stronger for prior audits. They primarily interpret this difference by claiming that the threat-of-audit letters only impact the probability of audit, while actual audits also change the probability of detection conditional on audit. [Advani, Elming and Shaw \(2019\)](#) exploit a random audit program to focus on the dynamic features of auditing and provide an alternative, yet complementary, interpretation of the mechanisms underlying the strength of deterrence effects: conducted audits update the information held by tax authorities and convey an auditee-specific threat that the latter is under scrutiny.

Insurance fraud has also been scrutinized in the academic literature. Building on the costly state verification⁵ framework developed by [Townsend \(1979\)](#), the theoretical literature studies insurance fraud through the scope of the optimal design of contracts and verification strategies under deterministic and random auditing (see [Gollier \(1987\)](#), [Huberman, Mayers and Smith \(1983\)](#), [Bond and Crocker \(1997\)](#), [Picard \(2000\)](#)). [Dionne, Giuliano and Picard \(2008\)](#) extend the situation to signal-based auditing in a Stackelberg model, i.e., with commitment. They show that optimal investigation strategies have a deterrence role, whereby an audit should still be conducted for some claims with negative expected auditing proceeds, underlying the importance of commitment. Solutions to the no-commitment case, initially analyzed by [Graetz, Reinganum and Wilde \(1986\)](#) and [Mookherjee and Png \(1989\)](#), are suggested in the insurance literature: [Picard \(1996\)](#) and [Boyer \(2000\)](#) mention resorting to a common agency to centralize insurance fraud investigations, while [Khalil \(1997\)](#) and [Boyer \(1999\)](#) focus on overcompensation of agents as a way to limit fraud. [Schiller \(2006\)](#) shows that the introduction of detection systems allowing the insurers to condition their audits on the information provided, in line with the recent developments in machine learning techniques for detection, renders auditing more effective and allows to reduce overcompensation. [Krawczyk \(2009\)](#) shows that the dynamic nature of the interaction between the insurer and the auditee might help alleviate the commitment

⁵An alternative to costly state verification is the costly state falsification framework introduced by [Crocker and Morgan \(1998\)](#).

problem and suggests that the insurer should spend resources on signaling their anti-fraud attempts. More recently, [Aboutajdine and Picard \(2019a\)](#) analyze the interaction between the roles of auditing as a deterrence and a learning mechanism in a dynamic setting, revealing a reputation-based deterrence effect: the threat faced by auditees of ruining one's reputation (if found to be defrauding) strengthens the deterrence effects of auditing. Empirically, the insurance literature has extensively focused on the statistical features of detection methods (see [Derrig \(2002\)](#), [Tennyson and Salsas-Forn \(2002\)](#), [Major and Riedinger \(2002\)](#), [Artís, Ayuso and Guillén \(2002\)](#), [Brockett et al. \(2002\)](#), [Viaene et al. \(2002\)](#)). Notably, [Tennyson and Salsas-Forn \(2002\)](#) discuss the role of auditing for both fraud detection and fraud deterrence, and show how auditing patterns are consistent with auditors taking into account both roles. [Pinquet, Ayuso and Guillén \(2007\)](#) rely on a bivariate probit model to estimate fraud rates when auditing is endogenous, i.e., when fraudulent cases are overrepresented in the sample and introduce selection bias. However, neither examine the auditees' reaction to investigation nor provide an estimation of the deterrence impact.

To the best of my knowledge, there has been no prior attempt at assessing the deterrence causal impact of insurance audits on fraud. This is especially striking when compared to the profusion of empirical investigations of deterrence in tax audits. One possible reason is that tax evasion has been at the heart of the public debate in the recent decade. Another plausible explanation consists in the scarcity of publicly available databases, as opposed to the case of tax evasion, and in the difficulty of accessing private ones: industrial data contains sensitive information about the firm's business processes, and it must comply with the local legislative constraints regulating the access and use of individual (health) data.⁶ This article is therefore a first step towards better understanding and quantifying the deterrent effects of insurance audits.

Methodologically, I rely on count data models, namely fixed-effects Poisson and NB1 negative binomial regression models, as introduced in [Hausman, Hall and Griliches \(1984\)](#). This is justified by the dependent variable being a number of submissions by a given optician at a given month, and having a large mass on zero. Count models have also notably been used to model doctor visits (e.g., [Cameron et al. \(1988\)](#)); here, we are interested in visits to an optician. Because the data is a short panel (i.e., small T), the question of possible inconsistency because of the incidental parameter problem⁷ must be addressed: while the Poisson model is immune to it ([Wooldridge \(1999\)](#)), the negative binomial's case is debatable ([Allison and Waterman \(2002\)](#)),

⁶See for example GDPR in the EU: <https://eugdpr.org/>.

⁷[Lancaster \(2000\)](#)

Greene (2004), Guimarães (2008)).

The rest of the paper is organized as follows. Section 1.2 explains the research context, describes the data and discusses my empirical strategy. Section 1.3 presents my empirical results. Section 1.4 discusses my results and suggests some extensions. Section 1.5 concludes.

1.2 Research Setting and Empirical Strategy

1.2.1 Research Setting

1.2.1.1 The French Health Insurance System

This research is conducted in the context of the French health insurance system, which has many specificities. Firstly, the insurance coverage is divided between two main types of entities: the (mandatory) public social security and the (optional) mutuals. Depending on the type of service, the shares of the price covered by each entity vary significantly. Secondly, most reimbursements are made in the context of a third-party payment system, called “Tiers-payant”.⁸ Service Providers⁹ (SPs) submit claims to an insurer in the name of their clients, the policyholders, and directly receive the reimbursement. Another reason SPs play a central role in the fraud process is that they act as certifiers, and a policyholder can almost never defraud without colluding with an SP. An SP can even defraud on his own, by submitting fictive claims. Therefore, SPs rather than policyholders are the focus of auditing in this setting. Finally, since a large share of health services are provided by the public sector, many health service providers have fewer incentives to defraud as their salary is fixed and not indexed on their economic activity. For instance, it is very rare for hospital doctors to defraud, and, when they do, their motive is usually to allow the patient to fall within the insurance reimbursement conditions.

This is not the case for opticians who act as privately owned businesses, and we focus on them hereafter. The case of opticians is of particular interest for several reasons. Firstly, social security only covers an insignificant share of the price of optical products, while the mutual reimburses most of the costs. This situation fosters fraud through two channels: it is less risky to defraud when only one entity, the mutual, has incentives to audit, and, because of the bad reputation private insurers have for nitpicking, defrauders may be less prone to suffer from remorse due to the moral cost of cheating.¹⁰ Secondly, opticians provide a tangible product (glasses, lenses) in addition to a service. This makes proving fraud easier, as there are receipts and invoices involved

⁸<https://www.service-public.fr/particuliers/vosdroits/F167>

⁹Doctors, pharmacists, opticians, ...

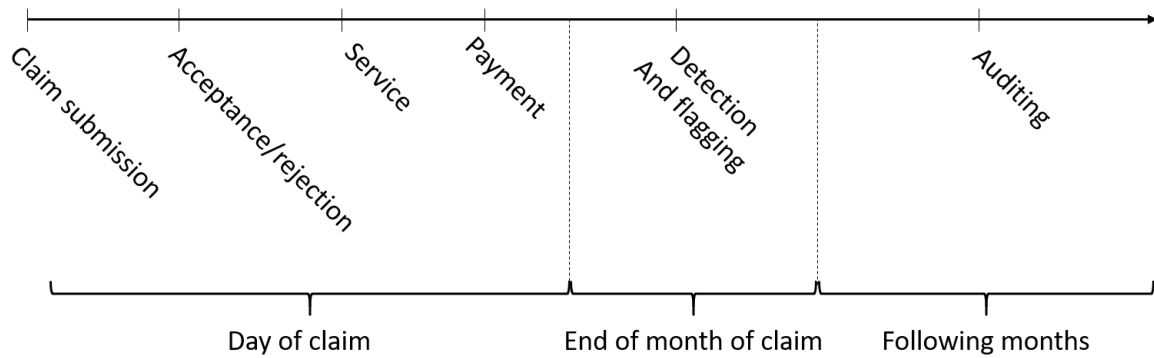
¹⁰See Tennyson (2008)

in the process since the optician himself must purchase the materials from manufacturers. The insurer can then ask for these supporting documents to check the validity of a claim. Finally, French opticians are interesting in that they may also work as audiologists. This allows me to control for a potential generality of the deterrence effects. I can then test if an audit targeting an SP for his audiology activity has some spillovers on his optical claim submission patterns.

1.2.1.2 Claiming, Detecting, Auditing and Deterring

This study takes place after the deployment of new detection systems developed by IBM France and the mutual PRO BTP. These systems help inform the mutual's decision to accept or refuse the submission of the claim *ex-ante*, and to audit the claim *ex-post*.

Figure 1.1 – TIMING OF CLAIM SUBMISSION AND AUDITING



Claiming The process starts with an optician submitting a coverage request electronically to the mutual in the name of one of its policyholders. The policyholder may be visiting the optician for legitimate reasons, or they may be colluding to disguise an illegitimate coverage request as a valid one. A common case involves colluders who request a reimbursement for regular glasses, when the reimbursed is actually used to cover the price of sunglasses which are not covered by the insurance contract.¹¹ Fraud schemes can be very creative and leverage clauses of the contract for other purposes than their original one. In some extreme cases, there may be no policyholder involved, when opticians submit a totally fictive coverage request to the mutual and cash in all the reimbursement.

The insurer receives the request and conducts a quick analysis to determine whether it falls

¹¹The sharing of the undue reimbursement between the policyholder and the optician may involve a bargaining process, but we do not include this aspect in our analysis. Note however that cases where the policyholder gets all the surplus are not rare, as the optician still benefits from using all of it to reduce the charged price for the consumer and encourage him to buy expensive glasses.

within the terms of the contract and to assess its legitimacy. This is just a preliminary superficial analysis and does not correspond to the actual thorough detection process. Then it either accepts or rejects the request. If it is accepted, it becomes a confirmed claim, the optician proceeds to deliver the service to the policyholder and the reimbursement is directly made to the optician.

Detecting At the end of the corresponding month, all accepted claims are thoroughly analyzed in the deployed detection system. The detection process involves several statistical methods such as rule-based red-flagging, anomaly detection and supervised learning. Its output is a flagging of claims depending on the likelihood of being fraudulent, and the results are transferred to an SIU within the mutual.

Auditing The SIU selects claims that it deems suspicious enough, based on the decision systems' outputs. The selection for auditing criteria are not fixed, and the criteria evolve in time and are based on discussions within the auditing teams. Therefore, while flagging is not random, it is hard to control for it explicitly without having access to an actual systematic selection process.

Detering Since the introduction of the new detection systems, the changes in the auditing process are twofold: first, because of the rationalization and automation of the process, it became possible to conduct more audits at lower costs. Second, thanks to the more sophisticated analytical process, the audits became more accurate and the probability of finding fraud conditional on an audit increased. According to the theoretical literature on criminal behavior and fraud ([Becker \(1968\)](#), [Allingham and Sandmo \(1972\)](#), [Dionne, Giuliano and Picard \(2008\)](#)), the equilibrium levels of fraud should decrease as a response to the increase in the probability of being investigated. However, this is under perfect information, and auditees should instead be expected to learn progressively about the shift in auditing patterns, through experiencing more frequent audits.

1.2.1.3 Data

My dataset includes a total of 260,592 observations: it encompasses all the claims submitted by 16,287 opticians to the mutual during the 16 months between August 2016 and November 2017. [Table 1.1](#) defines the main variables and [Table 1.2](#) presents some descriptive statistics.

Over the considered period, there have been 1,088,680 claims submitted to the mutual. [Figure 1.2](#) shows their distribution. The average optician submitted 4.18 claims per month. Standard deviation is 6.03, meaning claim count data is over-dispersed. Decomposing the variation in claims into between and within variation reveals heterogeneity with regards to the individual characteristics of opticians and to the timing of claims. First, the between standard deviation is

Table 1.1 – DEFINITION OF VARIABLES

Variables	Description
Claims	Number of claims submitted in a given month by a given optician
ClaimsH	Number of non suspicious claims ex-ante
ClaimsF	Number of suspicious claims ex-ante
Audits	Number of audits of any type
Warn	Number of warning letters, not specific to any given claim
DocR	Number of letters requesting supporting documents for specific claims
Rem	Number of letters reminding the optician to submit the supporting documents
Audio	Number of letters requesting supporting documents for specific audiology claims by the optician
Cxxx	Cumulative variable: e.g., CAudits indicates the cumulative number of received audits since the deployment of detection systems

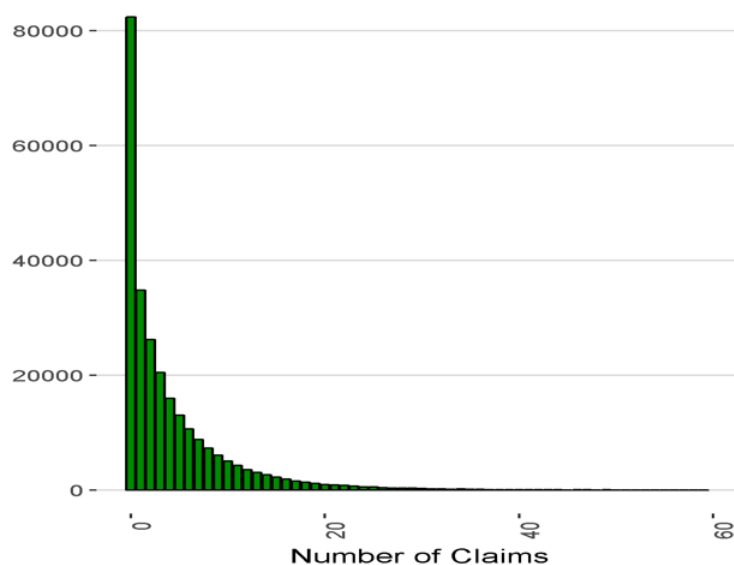
Table 1.2 – SUMMARY STATISTICS

Variables	Mean	Sd	Between Sd	Within Sd	Min	Max
Claims	4.18	6.03	5.18	3.08	0	101
ClaimsH	3.80	5.39	4.65	2.72	0	84
ClaimsF	0.37	1.56	0.90	1.27	0	97
Audits	0.0097	0.12	0.04	0.12	0	16
Warn	0.0027	0.06	0.02	0.05	0	4
DocR	0.0046	0.09	0.03	0.08	0	10
Rem	0.0013	0.04	0.01	0.04	0	6
Audio	0.0010	0.05	0.01	0.04	0	16
CAudits	0.26	1.03	0.99	0.28	0	32
CWarn	0.021	0.18	0.14	0.11	0	5
CDocR	0.12	0.50	0.46	0.19	0	18
CRem	0.05	0.37	0.36	0.08	0	15
CAudio	0.01	0.14	0.12	0.07	0	20
$N = 260, 592$	$N = i \times T$	$i = 16, 287$	$T = 16$			

5.18. This is due to heterogeneity in the claiming profile: some opticians submit large monthly numbers of claims, reaching a maximum of 101 claims, while others only seldom submit. This is a consequence of many factors: some opticians correspond to major nationwide chain stores while others can be rural small businesses. Some opticians may also interact more often with policyholders of the mutual than others. This heterogeneity in claiming intensity might also be a consequence of fraud: assuming honest claiming is independent of fraud, then defrauders would submit more claims than honest opticians on average. Second, the within standard deviation of 3.08 is a consequence to the seasonality of claiming patterns, as illustrated in Figure 1.3a. Opticians submit fewer claims during months with low economic activity, especially during August, as opposed to June and July. Submissions are also larger during particular periods such as December, which is partly related to policyholders trying to use their insurance before the

resetting of the contract for the new civil year.

Figure 1.2 – DISTRIBUTION OF CLAIMS

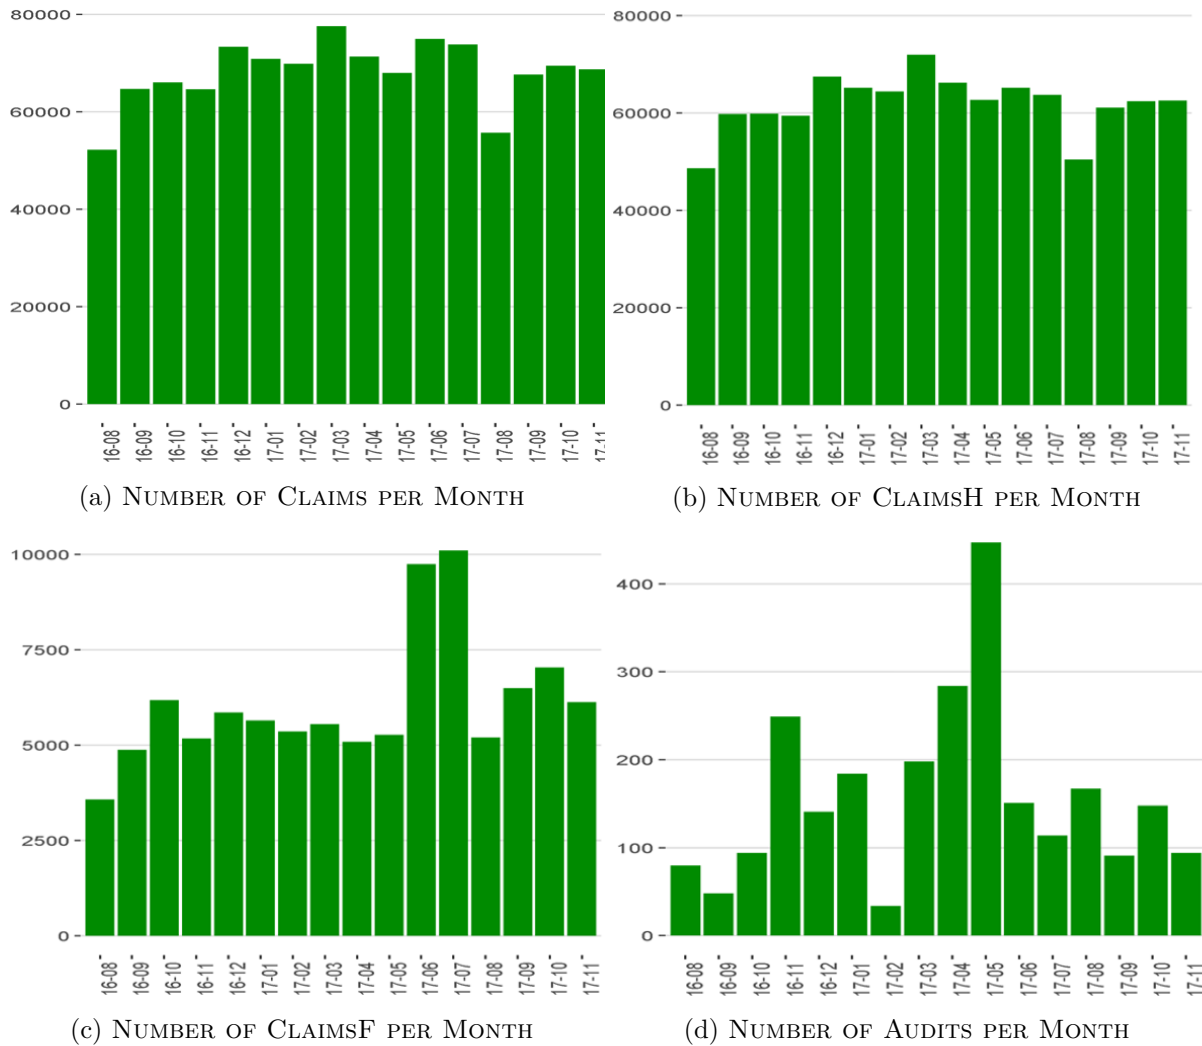


The volume of suspicious¹² claims remains marginal when compared to the volume of non suspicious ones: a monthly average of 3.80 claims per optician are believed to be honest, while a monthly average of 0.37 claims per optician are believed to be fraudulent. On average, the fraudulent claims represent 8.8% of the submitted claims, which is consistent with the 10% estimate commonly used in the industry. Both fraudulent and honest claim counts are overdispersed, in both the between and the within dimensions. However, while non suspicious claims display similar time patterns to the general claims (see Figure 1.3b), suspicious ones have a higher within standard deviation than their between standard deviation. This is because opportunistic fraud timing depends highly on the season. As illustrated in Figure 1.3c, the number of suspicious claims is remarkably larger during June and July. This is because an important share of fraud for vision insurance consists in purchasing sunglasses, which are seldom covered by the policy, while declaring regular glasses to the insurer, in order to reduce the former's price.

Regarding audits, a monthly average of 0.0097 claim per optician is audited. Auditing is therefore very rare, as it is costly and capacity constrained. It is also endogenous and targeted towards suspicious individuals, which explains the large standard deviation at 0.1256. A between standard variation of 0.04 shows that some opticians receive particularly intense scrutiny, and the maximum number of audits received by a given optician at a given month is as large as 16. The

¹²Based on the ex-ante suspicion scoring of detection systems.

Figure 1.3 – CLAIM SUBMISSION AND TIME DEPENDENCE



within variation is even larger with a 0.12 corresponding standard deviation, meaning auditing is also heterogeneous in timing. This is illustrated in Figure 1.3d. Understanding the concentration of auditing in some months rather than others is more subtle than the concentration of claims. It is an indirect function of the claiming pattern, as audit target claims in the months following submission, but it is also a function of unobservable random events such as a sudden change in the auditing policy of the mutual.

The audits can be of different types. Warning letters are basic deterrence tools that are sent after some suspicious general claiming patterns are detected, but without specific enough fraud signals at the claim level. If specific signals are observed, the SIU sends letters requesting supporting documents for the corresponding claims. These documents can be purchasing orders for glasses or lenses, or quotes for the offered service. The document requesting letters may also mention that the policyholder will be asked whether he did receive the service/product. The

latter is one of the cases where fraud is the less tedious to prove, especially when the optician defrauded on his own, without colluding with the policyholder. A third type, reminder letters, complementing the document request, may be sent in case of no response from the optician. This lack of response may for example be due to opticians not accepting to comply because they still do not perceive the threat as credible. Finally, an optician may receive a letter requesting documents for an audiology claim. It is not clear if such letter has an effect at the mutual level, where it would deter fraud even for the optical claims, or if it shifts fraud from audiology claims to optical ones.

Another quantity of interest is the cumulative number of incurred audits. In the context of the deployment of the new detection systems, the increase in the probability of being audited is only perceived by auditees through actually received audits. Thinking of auditees as Bayesians who update their belief accordingly, the higher the number of audits received in the past, the higher the updated probability of being audited, and the lower the probability of fraud. During the considered 16 months period, an average optician had received an average of 0.26 audits in the previous months.

1.2.2 Empirical Strategy

1.2.2.1 Evaluation Framework

Poisson estimator We evaluate the impact of the auditing history on the claiming submission pattern through a two-way fixed-effects Poisson model. We rely on the panel structure of the data to control for both individual and time effects. This estimation specifies the conditional mean as an exponential regression function, i.e., the Poisson parameter μ for $i \in \{1, \dots, 16287\}$ and $t \in \{1, \dots, 16\}$

$$\mu_{it} = \mathbb{E}[y_{it} | \mathbf{x}_{it}, \alpha_i] \quad (1.1)$$

$$= \alpha_i \exp(\mathbf{x}'_{it} \beta) \quad (1.2)$$

$$= \alpha_i \lambda_{it}. \quad (1.3)$$

The Fixed-Effects Poisson estimator is obtained from the regression

$$\begin{aligned} \mu_{it} &= \mathbb{E}\left[Claims_{it} | CAudits_{it}, \boldsymbol{\tau}_t, \alpha_i\right], \quad \boldsymbol{\tau}_t = \{\tau_{st} = \mathbb{1}\{s = t\} | s \in \{1, \dots, 16\}\} \\ &= \alpha_i \exp\left(\beta_{Aud} CAudits_{it} + \sum_{s=1}^{16} \beta_s \tau_{st} + \varepsilon_{it}\right), \end{aligned} \quad (1.4)$$

for auditing impacts regardless of type, and from the regression

$$\begin{aligned}\mu_{it} &= \mathbb{E}\left[Claims_{it} | CWarn_{it}, CDocR_{it}, CRem_{it}, CAudio_{it}, \boldsymbol{\tau}_t, \alpha_i\right] \\ &= \alpha_i \exp\left(\beta_W CWarn_{it} + \beta_D CDocR_{it} + \beta_R CRem_{it} + \beta_A CAudio_{it} + \sum_{s=1}^{16} \beta_s \tau_{st} + \nu_{it}\right)\end{aligned}\quad (1.5)$$

for type-dependent auditing impacts.

The individual specific effect α_i is not identified in this model because of the incidental parameter problem, as detailed in Section 1.2.2.2. Therefore, the model does not identify the marginal effect of auditing, given for regressor x_j by

$$ME_{itj} = \frac{\partial \mathbb{E}[y_{it} | \mathbf{x}_{it}, \alpha_i]}{\partial x_{itj}} = \alpha_i \exp(\mathbf{x}'_{it} \boldsymbol{\beta}) \beta_j = \beta_j \mathbb{E}[y_{it} | \mathbf{x}_{it}, \alpha_i].$$

Instead, results are interpretable as semi-elasticities. For a one unit change in a discrete variable x_{itj} , the proportionate change in $\mathbb{E}[y_{it} | \mathbf{x}_{it}, \alpha_i]$ is

$$\frac{\mathbb{E}[y_{it} | x_{itj} + 1, \mathbf{x}_{-jt}, \alpha_i] - \mathbb{E}[y_{it} | \mathbf{x}_{it}, \alpha_i]}{\mathbb{E}[y_{it} | \mathbf{x}_{it}, \alpha_i]} = \frac{\alpha_i \exp(\mathbf{x}'_{it} \boldsymbol{\beta}) (\exp(\beta_j) - 1)}{\alpha_i \exp(\mathbf{x}'_{it} \boldsymbol{\beta})} = \exp(\beta_j) - 1 \approx \beta_j,$$

for β_j small.

NB1 estimator The Negative Binomial NB1 estimator includes an overdispersion parameter ϕ_i in the conditional mean specification

$$\mu_{it} = \mathbb{E}[y_{it} | \mathbf{x}_{it}, \alpha_i, \phi_i] = \frac{\alpha_i \lambda_{it}}{\phi_i}.$$

Contrary to the Poisson case where the expectation and the variance are equal, the conditional variance is now a multiple of the mean

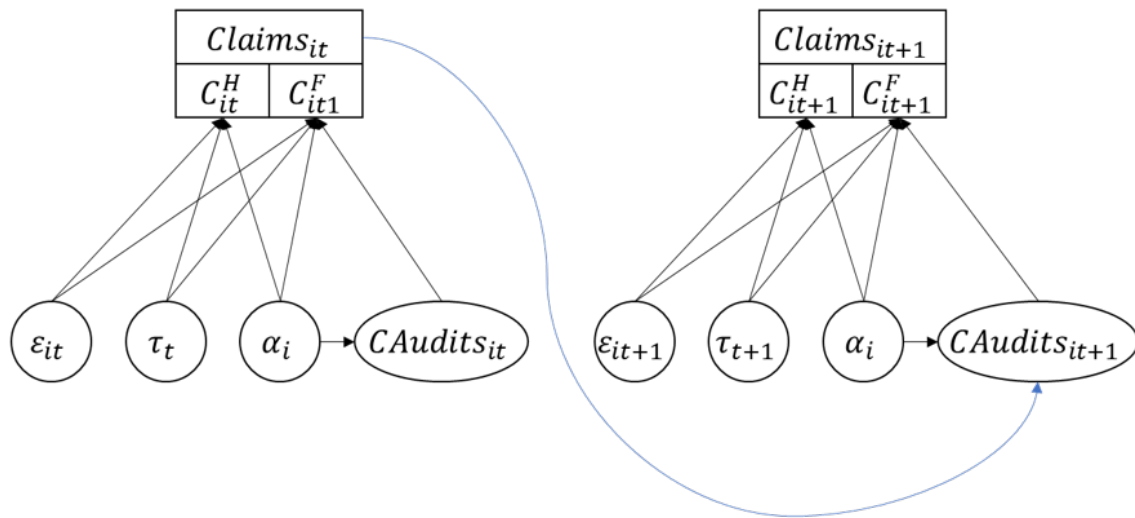
$$\mathbb{V}_{it} = \mathbb{E}[(y_{it} - \mu_{it})^2 | \mathbf{x}_{it}, \alpha_i, \phi_i] = \mu_{it} \times \left(1 + \frac{\alpha_i}{\phi_i}\right), \quad (1.6)$$

where the multiplicative term $(1 + \alpha_i/\phi_i)$ is larger than one, thus the overdispersion. The estimator is then obtained by conditional MLE from the conditional joint density for the i^{th} observation

$$\mathbb{P}\left[y_{i1}, \dots, y_{iT} \mid \sum_{t=1}^T y_{it}\right] = \left(\prod_t \frac{\Gamma(\lambda_{it} + y_{it})}{\Gamma(\lambda_{it}) \Gamma(y_{it} + 1)}\right) \times \frac{\Gamma(\sum_t \lambda_{it}) \Gamma(\sum_t y_{it} + 1)}{\Gamma(\sum_t \lambda_{it} + \sum_t y_{it})}. \quad (1.7)$$

1.2.2.2 Identification Strategy

Figure 1.4 – DIRECTED ACYCLIC GRAPH FOR EQUATION (1.4)



The two-way Fixed-Effects Poisson specification allows me to control for several confounders while avoiding the incidental parameter problem. I explain hereafter how I do so and discuss the type of effects the obtained coefficients identify. The Directed Acyclic Graph (DAG) in Figure 1.4 represents the relationship between variables.

Endogeneity of auditing and unobserved individual characteristics Auditing's main objective is to identify and target defrauders. It is therefore endogenous by nature. Running a Hausman test to decide between fixed and random effects yields a test statistic of 334.88, which is significant at the 0.01% level, thus confirming the severity of the selection bias in this setting.

Including individual fixed effects allows to control for a significant share of the unobserved variation in claiming generated by persistent individual characteristics. This is due to the fact that some of the main determinants of claiming, both honest and fraudulent, are fixed in time. Honest claiming is driven by location (urban vs. rural), marketing ability (attracting clients, building relationships), the type of optical store (major chain vs. personal business), etc. Fraudulent claim counts are influenced by morality and personal norms, technical ability (fraud is a complex process), social ability (convincing a visiting policyholder to collude and defraud), etc.

Incidental parameter problem

The fixed-effects Poisson case The major drawback of fixed-effects in non linear models for data from a short panel is the incidental parameter problem. The three leading exceptions are models with additive errors, multiplicative errors (e.g., Poisson), and the logit model. By analogy with differencing in the linear case, the α_i are eliminated in the Poisson case by quasi-differencing

$$\mathbb{E}[y_{it}|x_{i1}, \dots, x_{iT}, \alpha_i] = \alpha_i \exp(\mathbf{x}'_{it}\beta) = \alpha_i \lambda_{it}, \quad \text{under strict exogeneity.} \quad (1.8)$$

Averaging over time, $\mathbb{E}[\bar{y}_i|x_{i1}, \dots, x_{iT}] = \alpha_i \bar{\lambda}_i$. Then, subtracting from (1.8),

$$\mathbb{E}\left[(y_{it} - (\lambda_{it}/\bar{\lambda}_i)\bar{y}_i)|\alpha_i, x_{i1}, \dots, x_{iT}\right] = 0,$$

thus getting rid of the individual component α_i . A consequence, using the law of iterated expectations is

$$\mathbb{E}\left[x_{it}\left(y_{it} - (\lambda_{it}/\bar{\lambda}_i)\bar{y}_i\right)\right] = 0. \quad (1.9)$$

Finally, the fixed-effects Poisson estimator is obtained from the sample counterpart of equation (1.9).

The NB1 case The NB1 is also immune to the incidental parameter problem because, as shown in equation (1.7), the individual terms α_i and ϕ_i are not involved.

Endogenous timing of audits Audit timing is endogenous because, as illustrated in Figure 1.3d, audit numbers are not balanced in time. Explanations for this endogeneity include the fact that audits follow fraud-prone months, thus displaying seasonality, and are correlated with the general trend: audit numbers increase because of learning-by-doing at the SIU levels, but submissions also follow a positive trend (population growth, ageing of the population, see Asterès (2015) and Trevien (2015)).

The inclusion of time fixed effects deals with the endogeneity of the timing of audits.

Strict Exogeneity In equation (1.8), I assumed strict exogeneity to get rid of the non contemporaneous regressors. Strict endogeneity is usually written as

$$\mathbb{E}[\varepsilon_{it}|CAudits_{i1}, \dots, CAudits_{iT}, \boldsymbol{\tau}_t, \alpha_i] = 0. \quad (1.10)$$

I argue that strict exogeneity holds because of two reasons. Firstly, the impact of auditing, past and present, on claiming is entirely conveyed by the total number of incurred audits before the submission. Because my auditing explanatory variable is the cumulative number of audits $CAudits$, past audits do not directly affect current outcomes, and all relevant information is contained in $CAudits$. Secondly, the question of whether ε_{it} is independent of future audits ($CAudits_{is}|s > t$) must be addressed. Indeed, as shown by the blue arrow in Figure 1.4, the outcome $Claims_{it}$ from time t might affect future treatments. This is because an increase in claims with suspicious signals increase the probability of being selected for auditing in subsequent periods. My argument is that the generation of suspicious signals is already accounted for by the individual component α_i , the time component τ_{it} , and the contemporaneous audits $CAudits$. Indeed, fraud is very organized and planned, rather than opportunistic. If not enough collusion opportunities arise spontaneously, some opticians may call former customers, inviting them to visit the store and use their insurance for new glasses.

Auditing only affects the fraudulent subset of claims The DAG in Figure 1.4 only includes one arrow emerging from the $CAudits$ variable, directed towards the fraudulent component of the claim count. This assumption means that auditing only impacts fraudulent submissions. Indeed, honest claims stem from policyholders visiting an optician for a legitimate need, and there is no reason why an audited optician would refuse to serve legitimate policyholders. This assumption allows me to interpret any change in submission due to auditing as a change in fraudulent activity.

Average audit effects on the audited Finally, the regression estimates obtained from equations (1.4) and (1.5) can be interpreted as an average audit effect on the audited (AAEA), which is what I am interested in rather than the common average treatment effect (ATE). Indeed, my goal is to assess the efficiency of the deployed detection systems in terms of deterrence, so I am specifically interested in the deterrent effects conditional on the selection process generated by these systems. This AAEA should be expected to be higher than the ATE. An ATE estimate provided by a perfectly random selection would indeed include honest opticians for whom auditing impacts should be null, according to the previous assumption that audits should only affect fraudulent claims. These audits targeting honest claims would then drive the average estimated effect towards zero. My AAEA estimate, on the opposite, selects primarily defrauders for which there is deterrence potential. In addition, the selection process generated by the detection systems

is such that a defrauder is less likely to be audited as soon as he reduces fraud. Otherwise, had he continued being audited, the measured impact of auditing would have been lower, if not null.

1.3 Empirical Results

1.3.1 Impact of general auditing on claim submission

We consider first the general impact of past audits on submitted claims, regardless of the type of the audits. Table 1.3 provides the regression results for the regular Poisson estimator, the one-way Fixed-Effects Poisson (FEP) estimator, the two-way Fixed-Effects Poisson (2FEP) estimator, and the two-way Fixed-Effects NB1 (2FENB1) estimator.

Table 1.3 – GENERAL AUDITING IMPACT ON NUMBER OF SUBMITTED CLAIMS

Dependent Var Claims				
Indep Variables	Poisson	FEP	2FEP	2FENB1
Constant	1.430*** (0.011)			1.243*** (0.020)
CAudits	-0.002 (0.008)	-0.018** (0.008)	-0.030*** (0.008)	-0.042*** (0.010)
Sep16			0.214*** (0.009)	0.264*** (0.009)
Oct16			0.235*** (0.009)	0.284*** (0.009)
Nov16			0.214*** (0.009)	0.273*** (0.009)
Dec16			0.341*** (0.009)	0.389*** (0.009)
Jan17			0.307*** (0.009)	0.345*** (0.009)
Fev17			0.292*** (0.009)	0.342*** (0.009)
Mar17			0.398*** (0.009)	0.451*** (0.009)
Apr17			0.315*** (0.009)	0.362*** (0.010)
May17			0.268*** (0.009)	0.326*** (0.009)
Jun17			0.367*** (0.010)	0.424*** (0.010)
Jul17			0.353*** (0.010)	0.393*** (0.010)
Aug17			0.071*** (0.010)	0.098*** (0.010)
Sep17			0.266*** (0.010)	0.328*** (0.010)
Oct17			0.293*** (0.010)	0.353*** (0.010)
Nov17			0.283*** (0.010)	0.354*** (0.010)
N	260,592	258,656	258,656	258,656

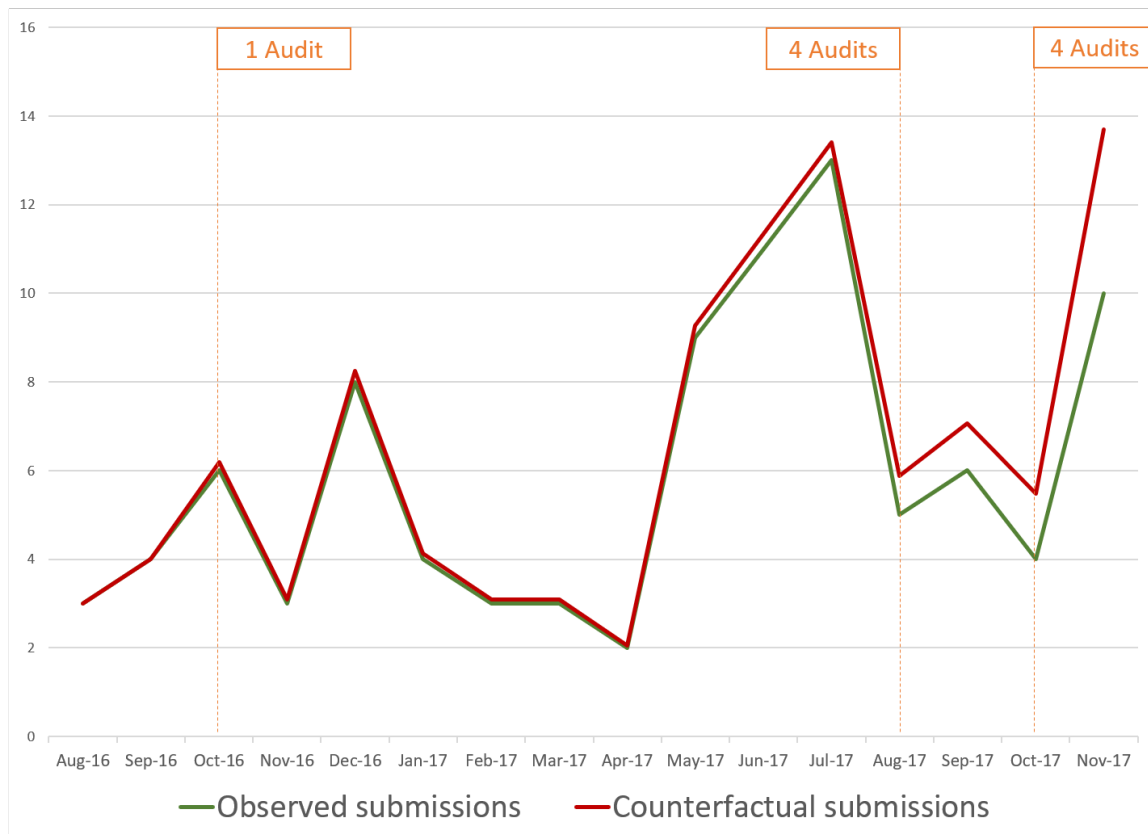
Standard errors in parenthesis. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

When unobserved heterogeneity is not accounted for, i.e., for the Poisson estimator, cumulative

past audits have no impact on claim submission. The corresponding coefficient is equal to -0.002 and not significant at any level. As expected from the discussion in Section 1.2.2.2, as soon as unobserved effects are accounted for through fixed-effects, results drastically change and become significant. Under the FEP specification, each additional audit reduces the average number of submissions by 1.8% and the result is significant at the 5% level.

Adding time fixed-effects in the 2FEP specification further reduces the coefficient, and each additional audit decreases claiming by 3.0%, a result significant at the 1% level. It confirms the endogeneity of the timing of audits and the existence of an omitted variable bias when not accounted for. The coefficients for the time dummies illustrate the seasonality and time trend of claiming. Note that all time coefficients are positive: since August 2016 serves as a basis, and August being the month with the least economic activity in France, it is intuitive that all other months witness higher submissions on average. In addition, even August 2017 has a positive coefficient, revealing the positive time trend in claim submission. This is confirmed by the fact that the other months that we observe twice in 2016 and 2017, that's to say September, October and November, have higher coefficients in 2017 than in 2016. Figure 1.5 illustrates the auditing deterrent effects by comparing observed and counterfactual submissions based on the 2FEP results.

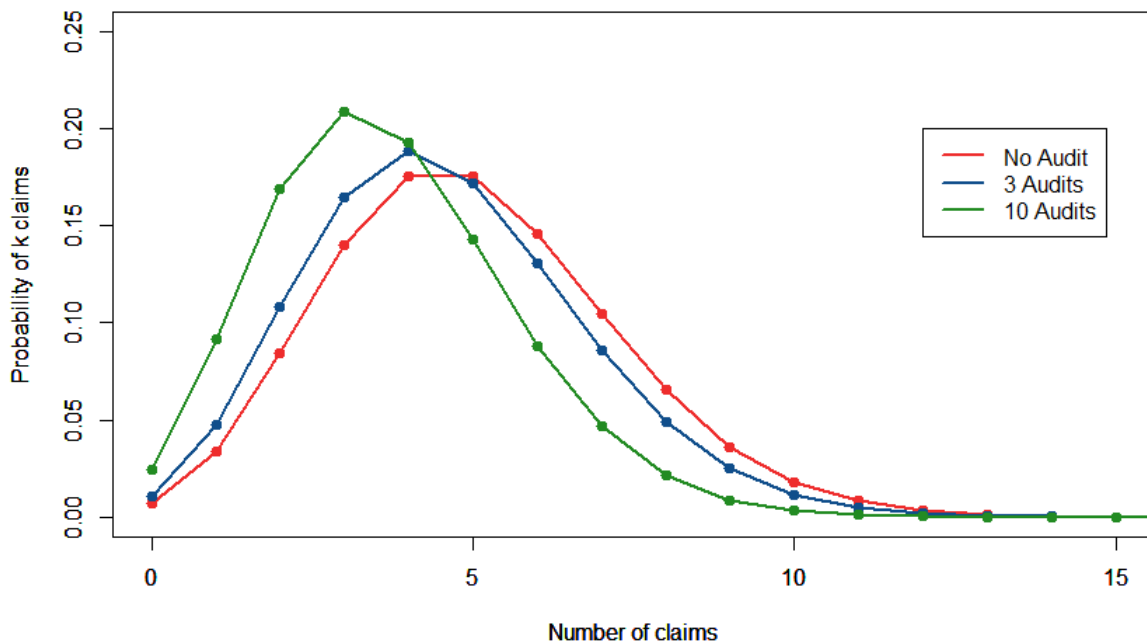
Figure 1.5 – COUNTERFACTUAL SUBMISSIONS AT THE INDIVIDUAL LEVEL



Finally, using the 2FENB1 model to account for over-dispersion confirms the negative impact of cumulative audits while all results remain significant at the 1% level. Each additional audit is now reducing submission counts by 4.2%. The 2FENB1 results are roughly similar to the 2FEP, although they indicate stronger auditing effects and stronger time effects.

These results underline the transition of equilibrium defrauding intensity from higher levels to lower ones, because of the change in the opticians belief about the probability of being audited, consistent with the predictions of crime and deterrence models (Becker (1968), Allingham and Sandmo (1972), Dionne, Giuliano and Picard (2008)). The impact of auditing on claim count can also be interpreted as a mass shift in the Poisson point mass function, as illustrated in Figure 1.6. Auditing shifts probability mass from submission numbers larger than the “no audits” Poisson parameter μ (= to 5¹³ in Figure 1.6) to smaller submission numbers. Moving from no audits to 3 audits, submission counts between 0 and 4 become more frequent, while submissions above 5 claims are less likely than before. Moving to 10 audits, the shift becomes more important for counts smaller or equal to 3, and there is almost no change in the probability of observing 4 claims, compared to the situation where 3 audits have been incurred.

Figure 1.6 – DETERRENCE AS A PROBABILITY POINT MASS SHIFT TO THE LEFT



¹³This is a fictive example to illustrate the change in the probability mass function. Because of the incidental parameter problem, we cannot estimate the individual specific parameter.

1.3.2 Impact on claim submissions depending on auditing type

I now estimate the effect of each type of audits on submission patterns separately. Regression results for each specification are shown in Table 1.4.

Starting again from the regular Poisson estimator results, I find that simple warning letter increase claim submissions by 1.9%, while being audited for audiology claims induces the optician to increase his optical submissions by 5.5%. Both results are significant at the 1% level. On the other hand, and while their coefficients are negative, neither document requesting letters nor reminder letters are significant at any level.

Table 1.4 – TYPE-SPECIFIC AUDITING IMPACT ON NUMBER OF SUBMITTED CLAIMS

Dependent Var				
Claims				
Indep Variables	Poisson	FEP	2FEP	2FENB1
Constant	1.425*** (0.010)			1.247*** (0.017)
CWarn	0.019*** (0.009)	0.009** (0.009)	-0.003 (0.009)	0.046*** (0.011)
CDocR	-0.020 (0.016)	-0.034** (0.016)	-0.042*** (0.015)	-0.062*** (0.020)
CRem	-0.035 (0.036)	-0.042 (0.040)	-0.061 (0.040)	-0.081** (0.037)
CAudio	0.055*** (0.015)	0.062*** (0.016)	0.041*** (0.016)	0.030 (0.028)
Sep16			0.213*** (0.009)	0.264*** (0.007)
Oct16			0.234*** (0.009)	0.284*** (0.009)
Nov16			0.213*** (0.009)	0.272*** (0.009)
Dec16			0.340*** (0.009)	0.389*** (0.010)
Jan17			0.307*** (0.009)	0.344*** (0.010)
Fev17			0.292*** (0.009)	0.342*** (0.008)
Mar17			0.397*** (0.009)	0.450*** (0.009)
Apr17			0.314*** (0.009)	0.361*** (0.009)
May17			0.267*** (0.009)	0.324*** (0.010)
Jun17			0.366*** (0.010)	0.421*** (0.009)
Jul17			0.351*** (0.010)	0.389*** (0.010)
Aug17			0.069*** (0.010)	0.094*** (0.010)
Sep17			0.264*** (0.010)	0.324*** (0.010)
Oct17			0.291*** (0.010)	0.348*** (0.012)
Nov17			0.281*** (0.010)	0.349*** (0.012)
N	260,592	258,656	258,656	258,656

Standard errors in parenthesis. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Once again, accounting for unobserved heterogeneity changes both the size and the significance of the coefficients. For the FEP estimator, document requesting audits are now significant at the 5% level and each one of them reduces claim submissions by 3.4% on average. Simple warning letters have a smaller positive impact which is less significant than before (5% level). Results do not change much for reminder letters and audio letters. When adding time fixed-effects in the 2FEP model, new estimates indicate no significant effect for warning letters, and document requesting letters reduce claim submissions by 4.2% (significant at the 1% level). The impact of reminder letters remains insignificant while the impact of audio letters remains significant at the 1% level, but decreases to 4.1%.

Finally, 2FENB1 results are this time different from the 2FEP. They refute any significant impact of audiology audits on optical claims. They also show that simple warnings have a significant positive impact (+4.6%) on claims submissions, while document requests have a stronger negative impact (-6.2%), both significant at the 1% level. Reminder letters are now significantly reducing claims submissions by 8.1% (significant at the 5% level).

This second set of results decomposes the impact of auditing depending on the type of audits, and shows the importance of commitment and credibility for deterrence. Because document requesting letters mention specific claims to be investigated, and because these claims are likely to be fraudulent, they convey an idea of credibility of the threat. On the contrary, warning letters are likely to be perceived as empty threats because they do not target specific claims: while they still show that the auditor is keeping an eye on submissions, these letters hint that it has no solid ground for litigation, and may even encourage the optician to defraud more. The case of reminder letters may be explained as follows: these letters are only sent after a document request letter, when the optician does not provide the documents in time or did not stop misbehaving. Then, they could capture a delayed deterrence reaction of opticians who were not sufficiently convinced by the initial letters. Finally, audio letters triggering an increase in optical submissions seems to indicate that auditees do not perceive the audit threat at the mutual level, but as specific to the activity targeted. It would indicate that when feeling monitored for audiology claims, defrauder opticians shift their illegitimate efforts to optical claims.

1.4 Discussion: Limits and Extensions

In this Section, I discuss the robustness of my results and suggest extensions. Section 1.4.1 examines the extent of issues that may arise from not correctly accounting for overdispersion and

from deviating from strict exogeneity. Section 1.4.2 suggests extensions that are implementable with the currently available data. Section 1.4.3 suggests other extensions that would require complementary data.

1.4.1 Limits

Robustness to overdispersion In order to account for the overdispersion of the claim data, I have included a two-way fixed-effects NB1 estimator in my results. While results from 2FEP and 2FENB1 were roughly similar in Section 1.3.1, they did differ to some extent in Section 1.3.2. I argue hereafter that 2FEP estimates should be preferred to 2FENB1 estimates.

First of all, notice that, suprisingly, the 2FENB1 estimates include a constant. More generally, time-invariant regressors are identified in this model, which lead Allison and Waterman (2002) to argue that it is not strictly speaking a fixed-effects model. In addition, Guimarães (2008) shows that the fixed-effects NB1 model places a very strong restriction on the relationship between the individual characteristic α_i and the NB1 overdispersion parameter ϕ_i .

One alternative is to estimate a regular NB2 model with fixed effects. However, this last model is not immune to the incidental parameter problem, although simulations by Allison and Waterman (2002) and Greene (2004) suggest that the resulting bias may not be too large for a moderately small panel dimension.¹⁴

On the other hand, the 2FEP estimator has very strong robustness properties as long as the conditional mean in equation (1.8) is correctly specified. The distribution of y_{it} given (\mathbf{x}_i, α_i) is unrestricted and there can be overdispersion or underdispersion in the latent variable model.¹⁵ Because of how large the cross section dimension of my data is ($i = 16,287$), consistency of the 2FEP estimator guarantees the robustness to overdispersion of the corresponding estimates.

Strict exogeneity Deviating from strict exogeneity would question the obtained estimated. However, I argue that it would not threaten the major interpretation of my results, that is that auditing has a negative effect.

If the strict exogeneity assumption is violated, then it is most likely because there is feedback from the explained variable to future explanatory variables. In other words, higher auditing in the future provides information about present error terms, in that positive error terms could be associated with more fraudulent claiming and thus with more future audits. Still, positive

¹⁴Greene (2004) considers moderately small to mean $T \geq 5$. My data therefore falls into this definition with $T = 16$.

¹⁵See Wooldridge (2002), Section 19.6.4 “Fixed Effects Poisson Estimation”.

correlation between current error and future audits would mean upward biased estimators. Therefore, in this case, a negative impact of auditing should be stronger than measured because of the bias.

Finally, future work could correct for lack of strict exogeneity by using the method of moments estimators proposed by [Wooldridge \(1997\)](#), to deal with lag models with possible feedback. He suggests residual functions using future dependent and explanatory variables that lead to conditional moment restrictions of the form

$$r_{it}(\beta) \equiv y_{it} - y_{it+1}[\exp(\mathbf{x}_{it}\beta)/\exp(\mathbf{x}_{it+1}\beta)], \quad t = 1, \dots, T - 1,$$

$$\mathbb{E}[r_{it}(\beta_0)|x_{i1}, \dots, x_{it}] = 0.$$

1.4.2 Extensions for the available dataset

Multi-valued treatments The models used in this analysis assume auditing impacts to be linear in the number of audits. One way to account for non-linearity is to resort to multi-valued treatment effects models described in [Cattaneo \(2018\)](#) and [Cattaneo \(2010\)](#).

Heterogenous treatments effect and timing In this article, I do not control for heterogeneity in timing and in treatment. Yet, as argued before, there must be heterogeneity in auditing effects, and audits are not delivered at the same time to different auditees. The recent literature on these topics for classic treatment settings (see [Goodman-Bacon \(2018\)](#) and [de Chaisemartin and D'Haultfoeuille \(2019\)](#)) has shown that consistent estimation is not straightforward.

Intensive margin and count-continuous models By focusing on auditing effects on claim counts, I only measure the extensive margin of auditing effects, and neglect the intensive margin. For example, some claims may still be submitted after audit, but with smaller monetary amounts. There is also a distinction to be made between the type of effect auditing has from the extensive and intensive margins perspective. The first case is when all subcomponents of a claim are illegitimate, and the claim would not be submitted under strong enough deterrence effects. The second case is when a claim with a legitimate basis is still submitted, but with a smaller corresponding amount, because the illegitimate component is abandoned due to deterrence effects. Finally the third case is a mixed one, where the claim has a legitimate basis, but the charged price, after reimbursement, is still too large for the policyholder. An illegitimate component could be added to the claim by the optician to encourage the policyholder to purchase the

product.¹⁶ Then, under strong enough deterrence, the claim would not be submitted even if it has a legitimate component.

One way to account for both the extensive and intensive margins is to use bivariate count-continuous models (see [Meghir and Robin \(1992\)](#) and [Praag and Vermeulen \(1993\)](#)).

Network effects and spillovers Direct auditing is not the only channel for deterrence effects. There may be spillovers through network effects that may spread the investigation threat more efficiently than expected. In the context of tax evasion, [Boning et al. \(2018\)](#) provide evidence of these network effects and control for connection through common tax preparers, geographic neighborhood, and parent-subsidiary relationship. This analysis can be extended to my dataset by defining connection between opticians as treating the same policyholder, or policyholders with a common ophthalmologist. This would specifically capture the effect of auditing on networks of defrauders, where opticians may collude between themselves, and eventually with an ophthalmologist, to exchange social security numbers of their policyholders and submit fictive claims to the mutual in a systematic and premeditated manner. Then, the first question is to understand whether auditing one of the members of this network has an impact on other members. The second question is to identify whether there are central agents in the networks that an optimal but capacity constrained auditing strategy should targeted first (see for example [Ballester, Calvó-Armengol and Zenou \(2006\)](#)).

1.4.3 Extensions with richer data

Post-audit outcomes to qualify claims as honest or fraudulent In the currently available dataset, distinction between suspicious and non-suspicious claims was made ex-ante, based on the preliminary detection systems' outputs. While this ex-ante qualification does give us information about suspicion and fraud levels, it is not perfectly correlated to the real status of the claim. For example, following a successful audit, future suspicion is systematically increased, independently of the signals generated by future claims. This artificially decreases the number of future non-suspicious claims and increases the number of future suspicious claims. Then the estimated impact of auditing on honest counts will be biased downwards while the estimated impact of fraudulent claims will be biased upwards. Post audit qualification would help circumvent this measurement issue.

¹⁶e.g., the policyholder needs glasses, but the price is still too large after reimbursement, so the optician declares lenses that he does not deliver, but instead uses the related reimbursement to reduce the price for the glasses.

Shift of fraud to competitors with less efficient counter fraud measures It would be interesting to control for the existence of outside options for defrauders, namely the existence of other insurers/mutuals with less sophisticated detection systems. The intuition is that a defrauder is more likely to be deterred by the mutual of interest if he can shift his fraudulent activity to a less audit prone insurer. But this would require data on competitors.

1.5 Conclusion

This paper provides evidence of the existence of deterrence effects in insurance audits and quantifies their impact on claim submissions. Using a two-way Fixed-Effects Poisson and NB1 estimators, I isolate these deterrence effects from other confounding factors such as unobserved heterogeneity and time distribution of audits.

I show that each additional audit incurred in the past reduces claim submission counts by 3%. In addition, specifically credible audits reduce submissions by 4.2% on average, while non-credible ones have no impact, or may even encourage fraud by conveying a lack of commitment from the auditor.

Yet, this is just a preliminary analysis of insurance auditing deterrence effects. There are several directions in which these results can be extended and complemented. Opening the black box of selection for audit is probably the most important one. Using signals gathered by an insurer to allocate her auditing resources would help instrument selection for audit and better control for the time varying component of endogeneity.

Chapter 2

Preliminary Investigations for Better Monitoring: Learning in Repeated Insurance Audits ¹

¹This chapter a joint work with Pierre Picard. We are grateful to seminar participants at CREST for helpful discussions and comments.

Abstract

Audit mechanisms frequently take place in the context of repeated relationships between auditor and auditee. This paper focuses attention on the insurance fraud problem in a setting where insurers repeatedly verify claims satisfied by service providers (e.g., affiliated car repairers or members of managed care networks). We highlight a learning bias that leads insurers to over-audit service providers at the beginning of their relationship. The paper builds a bridge between the literature on optimal audit in insurance and the exploitation/exploration trade-off in multi-armed bandit problems.

JEL Classification Numbers: D82, D83, G22, L21.

Keywords: Ex-post moral hazard, Insurance fraud, Optimal auditing, Learning.

2.1 Introduction

Claim fraud represents a serious threat to insurance markets: by artificially inflating the frequency and the cost of reported losses, defrauders lead to higher insurance premiums and they contribute to jeopardizing the efficiency of risk sharing mechanisms. Besides the free-riding problem that it poses, large scale fraud may even endanger the sustainability of the insurance markets that are prone to fraud.

Insurance claim fraud is sometimes referred to as a form of ex-post moral hazard in that it occurs after an (alleged) accident, for example when policyholders build up their claim or when they announce accidents that never actually happened. It essentially differs from ex-ante moral hazard through the associated timing (before/after the accident) and the modus operandi of addressing these inefficiencies: while principals tend to rely on a contract design approach to distort the agent's incentives in an ex-ante context (without being able to monitor the agent's effort), the ex-post situation is usually addressed through costly auditing in order to check what actually occurred.

The economic literature has mainly examined these issues through the lens of the *costly state verification* approach, whose foundations were laid by the seminal papers of [Townsend \(1979\)](#) and [Gale and Hellwig \(1985\)](#). Within this setting, it is assumed that the insurer can verify the true value of claims by incurring an audit cost.² The audit may be either deterministic, random or guided by signals perceived by the insurer. In particular, [Mookherjee and Png \(1989\)](#) establish that random auditing dominates deterministic auditing, while [Dionne, Giuliano and Picard \(2008\)](#) build a scoring methodology to show how audits are triggered by signals observed by the insurer. In one way or another, an optimal claim monitoring strategy achieves a trade-off between the additional costs of more frequent audits and the advantages of a more efficient fraud detection. The deterrence effect highlighted by [Dionne, Giuliano and Picard \(2008\)](#) is an example of such an advantage: they consider a setting where more frequent audits reduce the frequency of fraud, and they show that some individually unprofitable audits should be performed because of this deterrence effect.

Audit may also play an important role for gathering evidence about the auditee (e.g., does he seem to have a penchant for dishonest behavior?), an information that may be useful at later stages. Indeed, claimants (or service providers with whom they collude) may have some intrinsic and hard to observe characteristics that affect their propensity to defraud. Audit may help the

²[Bond and Crocker \(1997\)](#), develop the *costly state falsification* approach, where it is the defrauder who may incur some expenses to misrepresent her loss.

insurer to mitigate this informational asymmetry about claimants' type.³

The learning dimension is particularly relevant when it comes to repeated audits. Consider for instance the health insurance fraud case when there is a third party involved beside the insurer and the policyholders. Health service providers (doctors, opticians, pharmacists, etc.) play a central role since collusion between providers and policyholders is usually a necessary condition for fraud to take place. Furthermore, health care providers interact on a regular basis with the insurer, as they provide services to many policyholders and during several periods. Because of this repeated interaction, health insurers' anti-fraud efforts often focus on service providers as much as on policyholders. The same logic applies to property insurance, when insurers interact with car repairers or construction companies, sometimes within a network of affiliated service providers.

The purpose of the present paper is to investigate how such repeated interactions affect the optimal audit strategy. We will show that the insurer may find it optimal to perform unprofitable audits, because of a learning effect. In short, auditing is a way to gather information that can be used at later stages of the auditor-auditee interaction. This learning dimension may lead the insurer to perform audits beyond what would be optimal from a purely instantaneous standpoint.

To highlight this effect, we will rely on a simple model with two types of service providers (honest and dishonest) and where the likelihood of submitting an invalid claim is exogenous and type-dependent. Thus we abstain from analyzing the strategic interaction between the insurer and policyholders. We focus attention on the role of service providers as mandatory intermediaries who certify claims, without examining the collusion process between providers and claimants.⁴ Dishonest providers may certify invalid claims on purpose, while honest providers may only do it unintentionally. The insurer has beliefs about the type of each service provider and his decision consists in choosing the probability with which he audits each provider over the course of two consecutive periods. Auditing a claim allows the insurer to discover whether it is valid or invalid, but in the latter case, it does not reveal if the misbehavior was done on purpose or unintentionally. Auditing a claim allows the insurer to update his belief about the type of the service provider.

We find that, at the first period, the insurer has an incentive to perform some unprofitable audits, in order to improve his information about the service providers' type, and this additional

³Dionne, Giuliano and Picard (2008) introduce this hidden heterogeneity under the form of a cost reflecting the policyholder's moral sense that affects the probability of defrauding. Still, this cost remains unobservable by the insurer and does not come into play to assess the probability of a claim being fraudulent.

⁴Bourgeon, Picard and Pouyet (2008) analyze the collusion between service providers and policyholders when insurers have networks of affiliated providers.

information will allow him to more efficiently focus his auditing strategy at the second period. Ultimately, deviating from a strategy that would be guided by instantaneous expected gains, proves to be profitable. It corresponds to insisting on preliminary investigations early in the relationship to better monitor the agents later on.

This conclusion reveals an exploration/exploitation dilemma analogous to the multi-armed bandit problem in machine learning⁵. In this approach, a player is repeatedly facing a slot machine with multiple arms. He must choose an arm at each period, each arm providing a random payoff with imperfectly known stationary distribution. The player faces a trade-off between playing the most profitable arm according to early beliefs, and playing different arms in order to refine his beliefs. Exploring new arms induces an opportunity cost of not exploiting the arm that is the most profitable according to the current information, but this may allow the player to discover that other arms are in fact more profitable. Similarly, in our model, by trading current revenue for information, the better informed insurer gets a higher future payoff that compensates the initial loss.

The rest of the paper is organized as follows. In Section 2.2, as a preliminary stage, we consider a single period model where audit is not repeated and we characterize the corresponding optimal auditing strategy. In Section 2.3, we extend our model to account for repeated audits. We exhibit the competing roles of auditing as sources of revenue and information, and we define the insurer's dynamic optimization problem that will be solved by backward induction. Hence, we start by characterizing how available information is used at the second period and, in a second stage, we deduce how the first period audit should be performed. We show that the learning effect leads the insurer to audit more at the beginning of the relationship, with the magnitude depending on the informativeness of the audit and on the degree of short-sightedness of the insurer. In Sections 2.2 and 2.3, we restrain ourselves to a simple model where all claims have the same value. Section 2.4 extends our results to a more general setting with variable claim values. The final section concludes. Proofs are in the Appendix.

2.2 Single Period Auditing

2.2.1 Setting

Let us start by considering an insurer who interacts with a population of service providers (SPs) during a single period. SPs are mandatory intermediaries between insurer and insured.

⁵See Sutton and Barto (1998) and Bergemann and Välimäki (2008).

In particular, they certify the claims filed by policyholders, which means that they attest that the claims actually correspond to the value of the services paid by the policyholders following the event covered by the insurance policy. Each SP transmits exactly one claim with value normalized to 1 to the insurer, with the claim being either valid or invalid. Invalid claims should not lead to insurance payments. They may be transmitted either in good faith (for instance because the SP makes an error due to imperfect information about the circumstances of the loss) or in bad faith with the intention of defrauding.

SPs are heterogeneous when it comes to their propensity to transmit invalid claims. There are many possible determinants of this propensity, among which the sense of moral values, which is negatively correlated with the propensity to defraud, or the ability to build complex defrauding schemes.⁶ Hereafter, we consider that each SP may be either honest (H) or dishonest (D). Honest SPs only transmit invalid claims by error (they are always in good faith), while dishonest SPs may transmit invalid claims either by error or intentionally (they may be in bad faith). Hence a type H is less likely to transmit invalid claims than a type D. We include this aspect by defining probabilities $\mathbb{P}(Inv|H) = p_H$ and $\mathbb{P}(Inv|D) = p_D$ of submitting an invalid claim by type H and type D, respectively, such that $p_H < p_D$.

There is a continuum of SPs with mass 1 and the insurer has initial belief $\pi \in [0, 1]$ for each SP that represents the a priori probability that the SP is of type D. The prior π is distributed in $[0, 1]$ with density $f(\pi)$ and c.d.f $F(\pi)$ in the population of SPs. While we consider this prior as given, we may consider that it has been induced by signals (including the outcome of audits) that have been previously perceived by the insurer about each SP. These beliefs may be biased or not among SPs, i.e., the expected value $\int_0^1 \pi f(\pi) d\pi$ may or may not be equal to the true proportion of dishonest SPs in the continuum.

Each claim may be audited and, in that case, the insurer observes whether it is valid or invalid. The audit is costly and represents the fundamental constraint that the insurer faces when it comes to choosing audit targets. It costs c to investigate a claim, with $p_H < c < p_D$.

Claims found to be invalid are not paid, inducing a net proceed of $1 - c$. No penalty is paid to the insurer by SPs whose invalid claims are detected by the audit. Hence, auditing a claim is profitable (in expected terms) only when it has been certified by a dishonest SP. From the insurer's point of view, an SP with prior π transmits invalid claims with probability $\bar{p}(\pi) = p_D\pi + p_H(1 - \pi)$ and the corresponding expected net proceed of auditing is $\bar{p}(\pi) - c$.

⁶While defrauding in plain sight may occur (hoping for inattention of the insurer), it usually takes some effort to construct a defrauding scheme. For example, some opticians may provide sunglasses to their clients, but they certify that they have delivered regular glasses.

2.2.2 Auditing Strategy, Objective Function and Optimization Problem

For each SP, the insurer must decide whether an audit will be performed or not. We define an auditing strategy as a function $x(\cdot) : [0, 1] \rightarrow [0, 1]$ that assigns a probability $x(\pi)$ of being audited to each SP with belief π . For a given auditing strategy, the net expected proceed of audits is written as

$$\Omega(x(\cdot)) = \int_0^1 [\bar{p}(\pi) - c] x(\pi) f(\pi) d\pi. \quad (2.1)$$

The optimization problem of the insurer is written as

$$\begin{aligned} \max_{x(\cdot)} \quad & \Omega(x(\cdot)) \\ \text{s.t.} \quad & 0 \leq x(\pi) \leq 1 \quad \forall \pi \in [0, 1]. \end{aligned}$$

Lemma 2.1. *The single period optimal auditing strategy $x^*(\pi)$ consists in auditing all claims transmitted by SPs with associated beliefs $\pi \geq \pi^*$ and not auditing claims when $\pi < \pi^*$, i.e.,*

$$x^*(\pi) = \left\{ \begin{array}{ll} 1 & \text{if } \pi \in [\pi^*, 1] \\ 0 & \text{if } \pi \in [0, \pi^*) \end{array} \right\} = \mathbb{1}_{\{\pi \geq \pi^*\}},$$

where the threshold π^* is

$$\pi^* = \frac{c - p_H}{p_D - p_H} \quad \text{with} \quad \bar{p}(\pi^*) - c = 0.$$

Lemma 2.1 is unsurprising: one should only perform audits that are individually profitable, which amounts to focusing audits on SPs with π such that $\bar{p}(\pi) - c \geq 0$.

2.3 Two-Period Auditing: The Learning Effect

Because SPs take care of many policyholders, they repeatedly interact with the insurer. For the sake of simplicity, we assume that this interaction takes place during two consecutive periods $i = 0, 1$. There are different beliefs at the beginning of each period and the insurer's strategy is based on these beliefs. From now on, variables of interest will be indexed by the corresponding periods $(\pi_i, x_i, \Omega_i)_{i \in \{0,1\}}$. The insurer's inter-temporal objective function depends on both period specific objective functions Ω_0 and Ω_1 , the latter being weighted by $\gamma > 0$.⁷ His optimization

⁷If $\gamma \in (0, 1)$ it can be simply interpreted as a discount factor. Period 1 can also be viewed as the aggregation of all future proceeds without restriction about the value of γ .

problem is written as

$$\begin{aligned} & \max_{x_0(\cdot), x_1(\cdot)} \quad \Omega_0 + \gamma \mathbb{E}_0[\Omega_1] \\ & s.t. \quad 0 \leq x_0(\pi_0) \leq 1 \quad \forall \pi_0 \in [0, 1], \\ & s.t. \quad 0 \leq x_1(\pi_1) \leq 1 \quad \forall \pi_1 \in [0, 1], \end{aligned}$$

where \mathbb{E}_0 corresponds to the expected value operator at the beginning of period 0, i.e., before performing audit during this period.

2.3.1 Auditing as a Source of Information

Period 0 audits allow the insurer to update his belief at the beginning of period 1. Depending on whether an audit has been performed and, if it has been, whether the claim was valid or invalid (*Val* and *Inv*, respectively), the updated beliefs $\tilde{\pi}_1$ are deduced from initial beliefs π_0 through Bayes' Law:

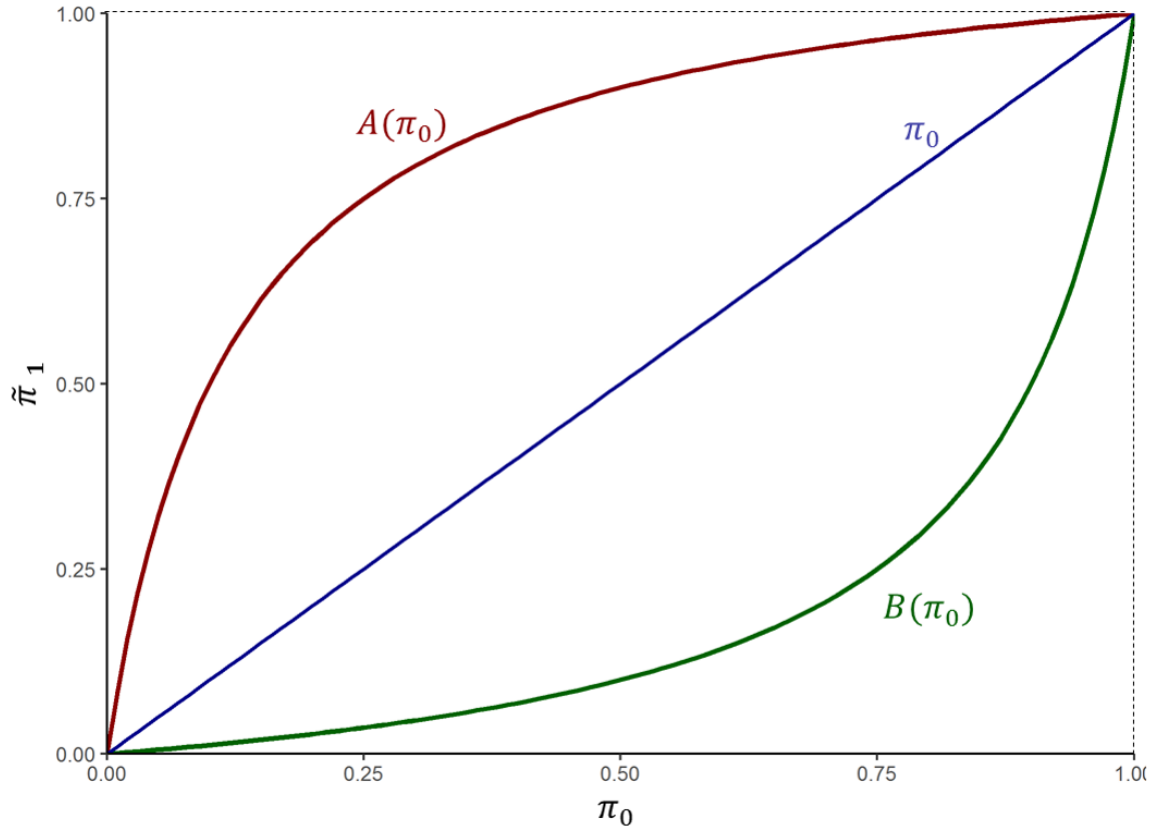
$$\tilde{\pi}_1 = \begin{cases} \mathbb{P}(D|audit, Inv) = A(\pi_0) = \frac{p_D \pi_0}{\bar{p}(\pi_0)}, \\ \mathbb{P}(D|audit, Val) = B(\pi_0) = \frac{(1-p_D)\pi_0}{1-\bar{p}(\pi_0)}, \\ \mathbb{P}(D|no\ audit) = \pi_0, \end{cases} \quad (2.2)$$

with

$$\begin{aligned} B(\pi_0) &< \pi_0 < A(\pi_0), \\ A' &> 0, \quad A'' < 0, \\ B' &> 0, \quad B'' > 0. \end{aligned}$$

Hence, $A(\pi_0)$ and $B(\pi_0)$ are the probabilities that the SP is dishonest if a period 0 audit revealed that the claim was invalid or valid, respectively. In particular, an invalid claim detected by audit leads the insurer to increase his beliefs that the SP is dishonest, i.e., $A(\pi_0) > \pi_0$, and it is the other way around if audit reveals that the claim was valid, i.e., $B(\pi_0) < \pi_0$. Of course, beliefs are unchanged if there is no audit.

For illustrative purposes, Figures 3.1 and 2.2 describe the degree of informativeness of an audit as a function of parameters p_H and p_D . In Figure 3.1, the graphs of functions $A(\cdot)$ and $B(\cdot)$ are symmetric on each side of the 45° line. This is due to the specific condition $p_H + p_D = 1$.

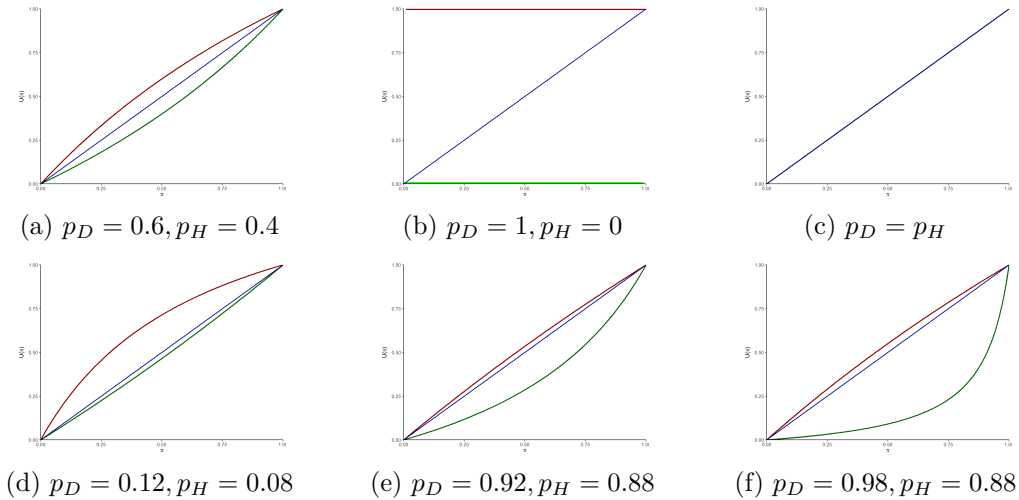
Figure 2.1 – UPDATING PRIORS ACCORDING TO THE AUDITING STATUS ($p_D = 0.9, p_H = 0.1$)


Maintaining this assumption, Figure 3.5a shows that bringing p_H and p_D closer makes both learning curves less concave/convex and closer to the 45° line, underlining the fact that the audit is less informative in this case. Figures 3.5b and 3.5c illustrate the extreme case where the audit is respectively totally informative ($p_H = 0$ and $p_D = 1$) and not informative at all ($p_H = p_D$, both types behave the same way). Relaxing the $p_H + p_D = 1$ assumption, Figures 3.5d, 3.5e and 2.2f exemplify the asymmetry of informativeness between invalidity and validity of a claim: in 3.5d, both probabilities of defrauding are rather low, so stumbling upon a valid claim does not say much, while finding a claim to be invalid induces a stronger change in the belief. The opposite happens in 3.5e where both types defraud often, with the validity status becoming more informative.

This aspect of auditing suggests some influence of period 0 auditing outcomes on period 1 auditing decisions. The information revealed at period 0 may lead to an expected efficiency gain at period 1. To express this idea, let us denote $\omega(\pi_1, x_1) = [\bar{p}(\pi_1) - c]x_1$ the expected gain of an audit performed at period 1 with probability x_1 under belief π_1 .

Let $\tilde{\pi}_1$ be the updated belief. This is a random variable defined by equation (3.1) whose distribution depends on initial beliefs π_0 and on the period 0 auditing probability $x_0(\pi_0)$.

Figure 2.2 – UPDATING FUNCTIONS FOR DIFFERENT PARAMETERS



Proposition 2.1. *The optimal period 1 optimal audit strategy $x_1^*(\cdot)$ is such that*

$$\mathbb{E}_0[\omega(\tilde{\pi}_1, x_1^*(\tilde{\pi}_1)) | \pi_0, x_0(\pi_0)] \geq \omega(\pi_0, x_1^*(\pi_0)) \quad \forall \pi_0 \in [0, 1]$$

with a strict inequality if there exists $\pi_1 \in [0, 1]$ such that $x_1^(\pi_1) \neq x_1^*(\pi_0)$ and $\mathbb{P}(\tilde{\pi}_1 = \pi_1 | \pi_0) > 0$.*

Proposition 2.1 implies that period 0 auditing increases the insurer's period 1 expected payoff if it affects the period 1 auditing strategy: its informational value translates into an increase in period 1 net proceeds, besides its period 0 income maximization value.

2.3.2 Inter-temporal Optimization Problem

Period 0 expected auditing profit is written as

$$\Omega_0(x_0(\cdot)) = \int_0^1 [\bar{p}(\pi_0) - c] x_0(\pi_0) f_0(\pi_0) d\pi_0,$$

where $f_0(\pi_0)$ is the density of prior beliefs.

The updating process corresponds to a mapping of period 0 beliefs into period 1 beliefs, thus changing the latter's distribution depending on the chosen period 0 strategy. Let $f_1(\cdot)$ be the density of period 1 updated beliefs. It depends on period 0 auditing strategy and then it is written as $f_1(\pi_1 | x_0(\cdot))$. Period 1 expected auditing profit is written as

$$\Omega_1(x_0(\cdot), x_1(\cdot)) = \int_0^1 [\bar{p}(\pi_1) - c] x_1(\pi_1) f_1(\pi_1 | x_0(\cdot)) d\pi_1.$$

The inter-temporal optimal auditing strategy of the insurer can now be characterized by

backward induction. At period 1, the insurer follows the optimal single period strategy $x_1^*(\cdot)$ characterized by Lemma 2.1. Hence the optimal period 1 expected net proceed of auditing $\Omega_1^*(x_0(\cdot)) = \Omega_1(x_0(\cdot), x_1^*(\cdot))$ is written as

$$\begin{aligned}\Omega_1^*(x_0(\cdot)) &= \int_0^1 [\bar{p}(\pi_1) - c] \mathbb{1}_{\{\pi \geq \pi^*\}} f_1(\pi_1 | x_0(\cdot)) d\pi_1 \\ &= \int_{\pi^*}^1 [\bar{p}(\pi_1) - c] f_1(\pi_1 | x_0(\cdot)) d\pi_1.\end{aligned}$$

Given this period 1 optimal strategy, the period 0 optimal strategy $x_0^*(\cdot)$ is an optimal solution to

$$\begin{aligned}\max_{x_0(\cdot)} \quad & \Omega_0(x_0(\cdot)) + \gamma \mathbb{E}_0[\Omega_1^*(x_0(\cdot))] \\ \text{s.t.} \quad & 0 \leq x_0(\pi_0) \leq 1 \quad \forall \pi_0 \in [0, 1].\end{aligned}$$

2.3.3 Effect of period 0 audit on the audit decision at period 1

In order to show how period 0 audit affects the decision to audit at period 1, let us define π_a and π_b by

$$\begin{aligned}A(\pi_a) &= \pi^*, & \pi_a &= \frac{\pi^* p_H}{\pi^* p_H + (1 - \pi^*) p_D}, \\ B(\pi_b) &= \pi^*, & \pi_b &= \frac{\pi^* (1 - p_H)}{\pi^* (1 - p_H) + (1 - \pi^*) (1 - p_D)}.\end{aligned}$$

One easily checks that

$$0 < \pi_a < \pi^* < \pi_b < 1.$$

We have $A(\pi) \geq \pi^*$ if and only if $\pi \geq \pi_a$ and $B(\pi) \geq \pi^*$ if and only if $\pi \geq \pi_b$. Hence, π_a is the lowest belief such that, if found invalid at period 0, the SP will be audited at period 1 and π_b is the highest belief such that, if found valid at period 0, the SP will be not be audited at period 1.

These thresholds lead us to Lemma 2.2 in which we express the probability of auditing at period 1 as a function of the period 0 belief and auditing outcomes.

Lemma 2.2. *The effect of period 0 audit on the audit decision x_1 at period 1 is characterized*

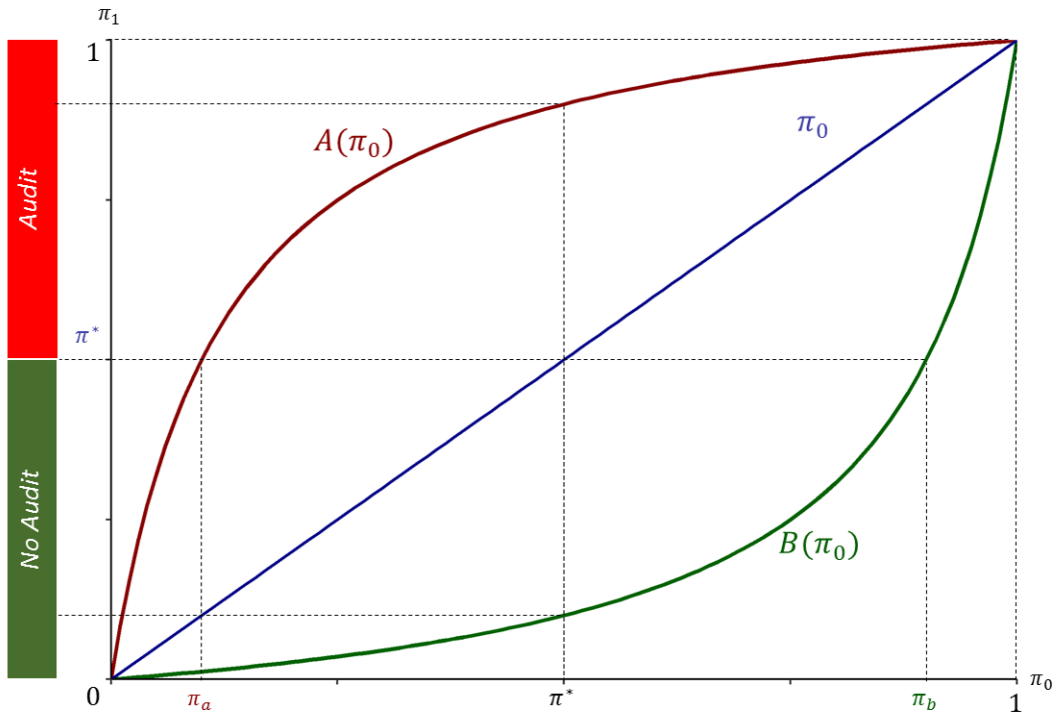
by Table 2.1.

Table 2.1 – PERIOD 1 AUDIT DECISION x_1 AS A FUNCTION OF PERIOD 0 BELIEFS AND AUDIT OUTCOMES

$\pi_0 \in$	$[0, \pi_a)$	$[\pi_a, \pi^*)$	$[\pi^*, \pi_b)$	$[\pi_b, 1]$
Audit and Valid	0	0	0	1
No Audit	0	0	1	1
Audit and Invalid	0	1	1	1

Figures 2.3 and 2.4 illustrate this relationship between the outcomes of the period 0 audit and the obtained posteriors.

Figure 2.3 – PERIOD 0 PRIORS AND PERIOD 1 AUDITING



Lemma 2.2 directly yields the probability of being audited at period 1 conditionally on π_0 and $x_0(\pi_0)$. This is written as

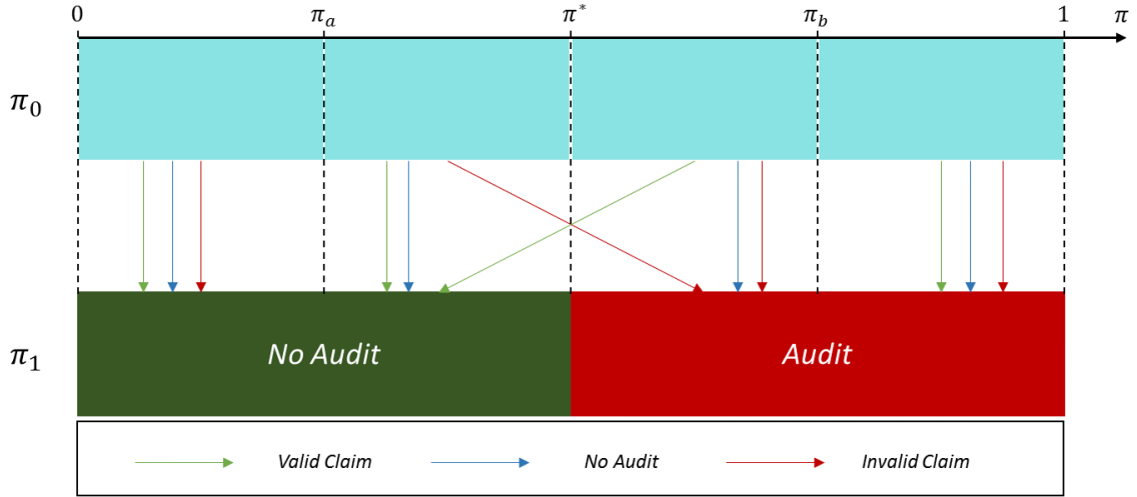
$$\mathbb{P}(\pi_1 \geq \pi^* | \pi_0 \in [0, \pi_a)) = 0,$$

$$\mathbb{P}(\pi_1 \geq \pi^* | \pi_0 \in [\pi_a, \pi^*)) = \bar{p}(\pi_0)x_0(\pi_0),$$

$$\mathbb{P}(\pi_1 \geq \pi^* | \pi_0 \in [\pi^*, \pi_b)) = \bar{p}(\pi_0)x_0(\pi_0) + 1 - x_0(\pi_0),$$

$$\mathbb{P}(\pi_1 \geq \pi^* | \pi_0 \in [\pi_b, 1]) = \bar{p}(\pi_0)x_0(\pi_0) + 1 - x_0(\pi_0) + (1 - \bar{p}(\pi_0))x_0(\pi_0) = 1.$$

Figure 2.4 – PERIOD 1 BELIEFS AS A CONSEQUENCE OF PERIOD 0 AUDIT OUTCOMES



For instance, when $\pi_0 \in [\pi^*, \pi_b)$, there will be an audit at period 1 either if, at period 0, an audit revealed an invalid claim or if there was no audit, which occurs with probability $\bar{p}(\pi_0)x_0(\pi_0)$ and $1 - x_0(\pi_0)$ respectively. The other cases can be interpreted similarly.

Lemma 2.3. *The inter-temporal objective function can be rewritten as*

$$\Omega_0(x_0(\cdot)) + \gamma \mathbb{E}_0[\Omega_1^*(x_0(\cdot))] = \int_0^1 [C(\pi_0) + K(\pi_0)x_0(\pi_0)] f_0(\pi_0) d\pi_0,$$

where functions $C(\cdot)$ and $H(\cdot)$ are defined in Table 2.2 and $K(\pi_0) = \bar{p}(\pi_0) - c + \gamma H(\pi_0)$. In particular, $K(\pi_0)$ is a continuous piecewise linear function such that

$$\begin{aligned} K(\pi_0) &< 0 \text{ if } \pi_0 \leq \pi_a \text{ and } K(\pi_0) > 0 \text{ if } \pi_0 > \pi^*, \\ K(\pi_0) &= \bar{p}(\pi_0) - c \quad \text{for } \pi_0 \in (0, \pi_a) \cup (\pi_b, 1), \\ K(\pi_0) &> \bar{p}(\pi_0) - c \quad \text{for } \pi_0 \in (\pi_a, \pi_b). \end{aligned}$$

Lemma 2.3 decomposes the inter-temporal objective function into components that explicit the impact of $x_0(\cdot)$ on current and future audit proceeds. $C(\pi_0)$ represents the proceeds of period 1 audits when the SP is not audited at period 0, and thus this term is not affected by $x_0(\pi_0)$. Beliefs π_0 in $[0, \pi_a)$ and $[\pi_a, \pi^*)$ are smaller than π^* , and thus, if the corresponding SPs are not audited at period 0, neither will they be at period 1. Hence, any benefit/loss from these beliefs is

Table 2.2 – DEFINITION OF $C(\cdot)$ AND $H(\cdot)$.

π_0	$C(\pi_0)$	$H(\pi_0)$
$[0, \pi_a)$	0	0
$[\pi_a, \pi^*)$	0	$p_D^2 \pi_0 + p_H^2 (1 - \pi_0) - \bar{p}(\pi_0)c$
$[\pi^*, \pi_b)$	$\gamma [\bar{p}(\pi_0) - c]$	$p_D^2 \pi_0 + p_H^2 (1 - \pi_0) - \bar{p}(\pi_0)c - [\bar{p}(\pi_0) - c]$
$[\pi_b, 1]$	$\gamma [\bar{p}(\pi_0) - c]$	0

necessarily the consequence of being audited at period 0, which gives $C(\pi_0) = 0$. This is different for beliefs in $[\pi^*, \pi_b)$ and $[\pi_b, 1]$ because they correspond to cases where initial beliefs π_0 are higher than π^* . Consequently, the SPs will be audited at period 1 and there are some proceeds from period 1 that do not depend on $x_0(\cdot)$, hence the presence of γ in $C(\cdot)$.

$K(\pi_0)$ represents the proceeds over the two periods that are affected by period 0 audit. $\bar{p}(\pi_0) - c$ in $K(\pi_0)$ represents the period 0 proceeds resulting from being audited at that period, while $H(\pi_0)$ corresponds to the period 1 proceeds that are affected by period 0 audits through the belief updating process. For instance, an SP with $\pi_0 \in [\pi_a, \pi^*)$ will be audited at period 1 if an audit revealed an invalid claim at period 0 (because $\pi_1 \geq \pi^*$ in that case) and the term $p_D^2 \pi_0 + p_H^2 (1 - \pi_0) - \bar{p}(\pi_0)c$ corresponds to the expected net proceeds. $H(\pi_0) = 0$ in $[0, \pi_a)$ and $[\pi_b, 1]$, although for different reasons: in $[0, \pi_a)$, regardless of the outcome of the audit, updated beliefs will remain below π^* and will never be audited at period 1, while in $[\pi_b, 1]$, whatever happens at period 0, all beliefs remain above π^* and the claim will always be audited at period 1. Function $K(\pi_0)$ is illustrated in Figure 2.5.

2.3.4 Inter-temporal Optimal Auditing Strategy

Proposition 2.2. *An optimal period 0 strategy $x_0^*(\cdot)$ is characterized by $\pi^{**} \in (\pi_a, \pi^*)$ such that*

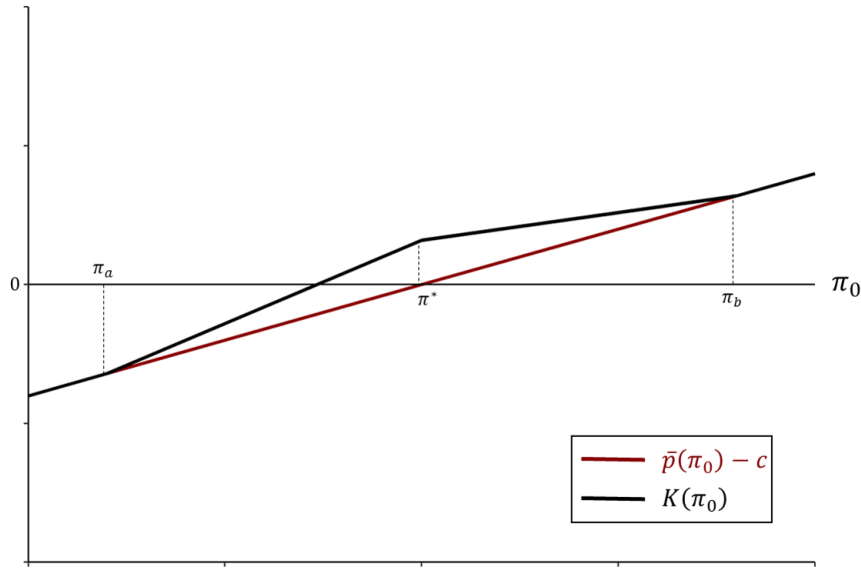
$$x_0^*(\pi_0) = 0 \text{ if } \pi_0 \in [0, \pi^{**}),$$

$$x_0^*(\pi_0) \in [0, 1] \text{ if } \pi_0 = \pi^{**},$$

$$x_0^*(\pi_0) = 1 \text{ if } \pi_0 \in (\pi^{**}, 1].$$

The threshold π^{**} is given by

$$\pi^{**} = \pi^* \frac{1 + \gamma p_H}{1 + \gamma p_H + \gamma(p_D - c)} < \pi^*,$$

Figure 2.5 – $K(\pi_0)$ FOR $(p_D = 0.9, p_H = 0.1, c = 0.5, \gamma = 1)$


and is such that

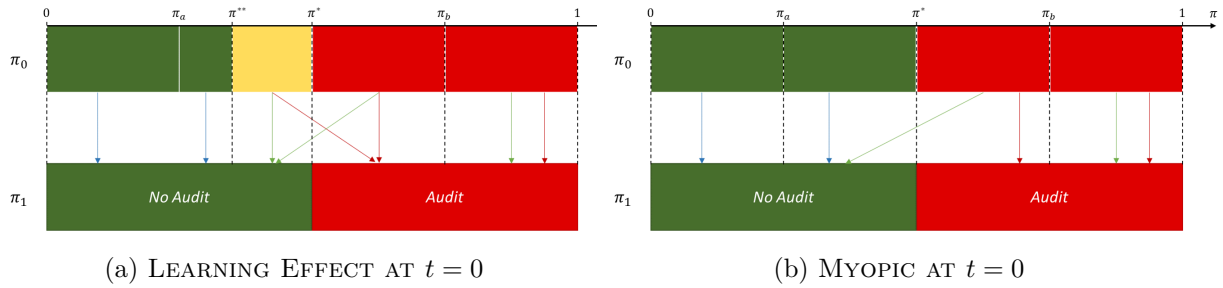
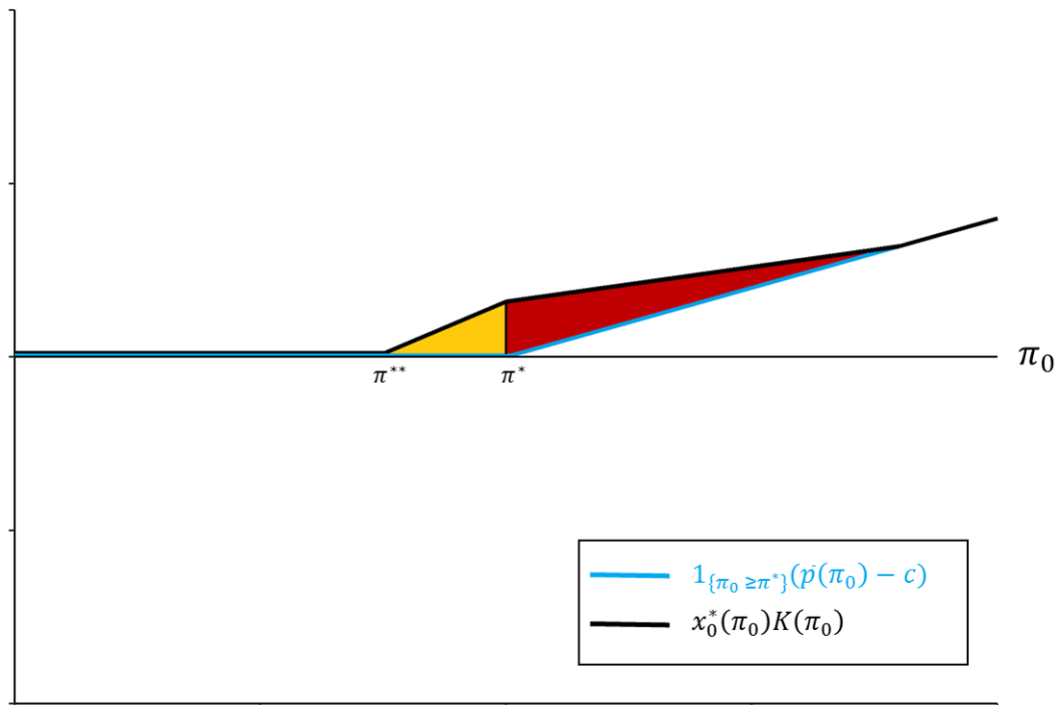
$$K(\pi^{**}) = 0 \quad \text{and} \quad \bar{p}(\pi_0) - c < 0 \quad \forall \pi_0 \in [\pi^{**}, \pi^*]. \quad (2.3)$$

Proposition 2.2 shows that, accounting for the impact of period 0 auditing on period 1, posterior beliefs lead the insurer to audit a higher number of SPs at period 0 than in the instantaneous audit problem analyzed in Section 2.2. The belief threshold above which claims should be audited is now π^{**} instead π^* , with $\pi^{**} < \pi^*$. An interesting aspect is that these additional auditees $\pi_0 \in [\pi^{**}, \pi^*)$ are such that the corresponding individual expect net proceeds of audit are negative (equation (2.3)): in other words, in spite of the negative impact on period 0 audit proceeds, the information gathered generates enough (discounted) profit at period 1 to compensate this initial loss. Figure 2.6 illustrates this deviation from the single period myopic auditing and Figure 2.7 shows in orange the additional period 1 positive net proceeds $\gamma \int H(\pi_0) f_0(\pi_0) d\pi_0$ that come from auditing down to π^{**} .⁸

The extent of the informational value of auditing is illustrated by the comparative statics properties of π^{**} . If $\gamma = 0$, i.e., if the insurer at time 0 does not care about period 1, then $\pi^{**} = \pi^*$ since the informational value of auditing at $t = 0$ serves no purpose. If $\gamma \rightarrow +\infty$, i.e., if the insurer only cares about period 1 profit, then $\pi^{**} \rightarrow \pi_a$ and he seeks to get the maximum information from period 0. Of course, there's no point in having π^{**} lower than π_a

⁸This result shows some similarity with the analysis of the deterrence effect by Dionne, Giuliano and Picard (2008). They show that some claim should be audited although the corresponding expected gain is negative. This is due to the deterrence effect of auditing: more intense monitoring discourages fraud and it should lead the insurer to audit below the individual claim profitability threshold.

Figure 2.6 – LEARNING VS MYOPIC


 Figure 2.7 – $x_0(\pi_0)K(\pi_0)$ FOR ($p_D = 0.9, p_H = 0.1, c = 0.5, \gamma = 1$)


as the additional information would not be useful at period 1. The new auditing limit, like the myopic one, also depends on c : if $c = p_H$ then $\pi^{**} = \pi^* = 0$ and if $c = p_D$ then $\pi^{**} = \pi^* = 1$, as in both case there is no more trade-off between information and revenue. Finally, if we write $c = \alpha p_D + (1 - \alpha)p_H$ with $\alpha \in (0, 1)$, then $\pi^* = \alpha$ and $\pi^{**} \rightarrow \pi^*$ when $p_D \rightarrow p_H$. When p_D comes closer to p_H , the separating power of the audit decreases until it becomes uninformative, and at the limit $p_D = p_H$ there is no more distinction between types D and H.

2.4 Variable Claim Value

A large part of the economic analysis of insurance fraud has focused attention on optimal auditing strategies when policyholders may file smaller or larger claims, and on the way the

insurer's audit strategy depends on the size of the claim.⁹ Let us consider how this approach may be affected by the learning mechanism.

2.4.1 Deterministic Value

As a preliminary step, let us consider the case where the size of all claims takes some arbitrary value $\ell \in \mathbb{R}^+$. This size is still fixed, but not necessarily equal to 1. The belief threshold π^* now depends on ℓ and is defined by

$$\begin{aligned} \pi^* &= 1 && \text{if } \ell \leq \frac{c}{p_D}, \\ \bar{p}(\pi^*)\ell - c &= 0 && \text{if } \frac{c}{p_D} < \ell < \frac{c}{p_H}, \\ \pi^* &= 0 && \text{if } \frac{c}{p_H} \leq \ell. \end{aligned}$$

Therefore

$$\pi^*(\ell) = \max \left\{ 0, \min \left\{ 1, \frac{c/\ell - p_H}{p_D - p_H} \right\} \right\}.$$

Equivalently we may define a threshold ℓ^* for the claim size

$$\ell^*(\pi) = \frac{c}{p_H + (p_D - p_H)\pi} = \frac{c}{\bar{p}(\pi)} \quad \forall \pi \in [0, 1],$$

and, for beliefs π , auditing is profitable if $\ell \geq \ell^*(\pi)$.

A straightforward extension of the results of Section 2.3, with the same claim size ℓ at each period, shows that the first period optimal auditing threshold becomes

$$\pi^{**}(\ell) = \pi^*(\ell) \times \frac{1 + \gamma p_H}{1 + \gamma p_H + \gamma(p_D - c/\ell)}.$$

As $\pi^{**}(\cdot)$ is strictly decreasing from $(\frac{c}{p_D}, \frac{c}{p_H})$ to $(0, 1)$, we can also define a function $\ell^{**}(\cdot) : [0, 1] \rightarrow [\frac{c}{p_D}, \frac{c}{p_H}]$ as

$$\begin{aligned} \ell^{**}(0) &= \frac{c}{p_H}, \\ \ell^{**}(\pi) &= (\pi^{**})^{-1}(\pi) \text{ for } \pi \in (0, 1), \\ \ell^{**}(1) &= \frac{c}{p_D}. \end{aligned}$$

⁹See Picard (2013) for a survey.

In the two period setting, a claim ℓ certified by an SP with associated belief π_0 will be audited if $\ell \geq \ell^{**}(\pi_0)$.

The set of claims Δ^* for which audit is profitable within a single period setting (with belief π and claim size ℓ) is defined by

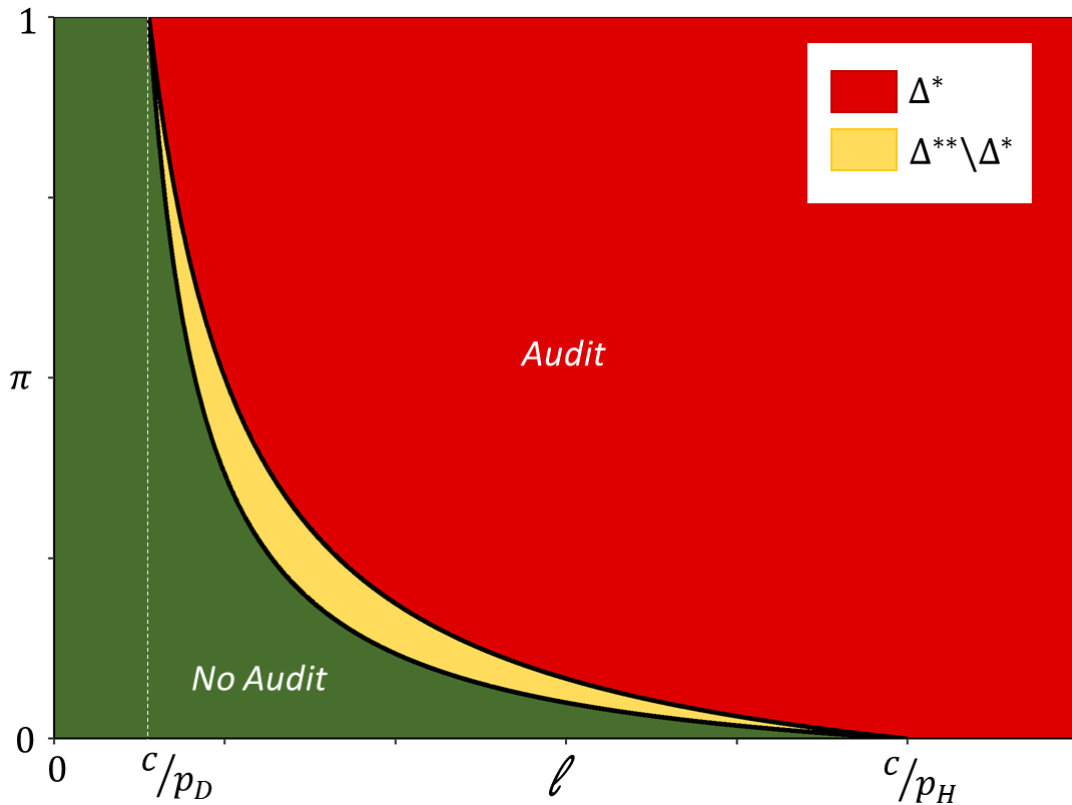
$$\Delta^* = \{(\pi, \ell) | \pi^*(\ell) \leq \pi \leq 1\} = \{(\pi, \ell) | \ell^*(\pi) \leq \ell\}.$$

In a two period setting, with claim size ℓ at both periods, claims should be audited at period 0 if $(\pi_0, \ell) \in \Delta^{**}$, where

$$\Delta^{**} = \{(\pi_0, \ell) | \pi^{**}(\ell) \leq \pi_0 \leq 1\} = \{(\pi_0, \ell) | \ell^{**}(\pi_0) \leq \ell\},$$

with $\Delta^* \subset \Delta^{**}$.

Figure 2.8 – AUDITING THRESHOLDS $\pi^*(\ell)$ AND $\pi^{**}(\ell)$ WITH A VARIABLE CLAIM VALUE ℓ



2.4.2 Random Homogeneous Value

Let us move on now to the more interesting case where the size of the claims is a random variable drawn from a known distribution at the beginning of each period. Let $\tilde{\ell}_i$ denote this

random variable with density $g_i(\cdot)$ and c.d.f. $G_i(\cdot)$ on $[0, L_i]$ for $i = 0, 1$. For simplicity, we assume that $\tilde{\ell}_i$ has the same probability distribution for both types of SPs, and thus $\tilde{\ell}_i$ and π_0 are independently distributed. The value of a claim is observed at each period before deciding to audit or not. Now, an audit strategy is written as $x_i(\pi_i, \ell_i)$ at period $i = 0, 1$.

Lemma 2.4. *The inter-temporal objective function can be written as*

$$\begin{aligned} \Omega_0(x_0(\cdot, \cdot)) + \gamma \mathbb{E}_0[\Omega_1^*(x_0(\cdot, \cdot))] \\ = \int_{[0,1]} \int_{[0,L_0]} [C(\pi_0) + K(\pi_0, \ell_0)x_0(\pi_0, \ell_0)] g_0(\ell_0) f_0(\pi_0) d\ell_0 d\pi_0, \end{aligned} \quad (2.4)$$

where

$$\begin{aligned} K(\pi_0, \ell_0) &= \bar{p}(\pi_0)\ell_0 - c + \gamma H(\pi_0), \\ H(\pi_0) &= \bar{p}(\pi_0)\Phi(A(\pi_0)) + (1 - \bar{p}(\pi_0))\Phi(B(\pi_0)) - \Phi(\pi_0), \\ C(\pi_0) &= \gamma\Phi(\pi_0), \end{aligned}$$

with $H(\pi_0) > 0$ if $\pi_0 \in (0, 1)$ and $H(0) = H(1) = 0$, and

$$\Phi(\pi) = \int_{\ell^*(\pi)}^{L_1} [\bar{p}(\pi)\ell_1 - c] g_1(\ell_1) d\ell_1,$$

where $\ell^*(\pi) = \frac{c}{\bar{p}(\pi)}$.

Lemma 2.3 and Lemma 2.4 are similar and can be interpreted the same way. In particular, the two terms in the integral of formula (2.4) correspond to the parts of cumulated expected proceeds according to whether they are affected by period 0 audit or not .

Proposition 2.3. *The optimal period 0 auditing strategy $x_0^*(\pi_0, \ell_0)$ is such that*

$$x_0^*(\pi_0, \ell_0) = 1 \text{ if } \pi_0 > \pi^{**}(\ell_0),$$

where $\pi^{**}(\cdot) : [0, L_0] \rightarrow [0, 1]$, with $\pi^{**}(\ell_0) < \pi^*(\ell_0)$ for all $\ell_0 \in [0, L_0]$.¹⁰

Proposition 2.3 extends Proposition 2.2 to the case of claims with variable size, and its interpretation is similar. In an instantaneous setting, where learning effects would be ignored, an audit should be performed if the belief π_0 is larger than $\pi^*(\pi_0)$. The threshold is decreasing

¹⁰Proposition 2.3 only states that $\pi^{**}(\ell_0) < \pi^*(\ell_0)$ for all ℓ_0 . Additional assumptions would allow us to show that $K(\pi_0)$ is monotonous and thus that claims should not be audited when $\pi_0 < \pi^{**}(\ell_0)$.

with ℓ_0 because the larger the claim, the larger the potential gain from auditing. When learning effects are taken into account, the threshold decreases. Claim (π_0, ℓ_0) should be audited at period 0 when $\pi^{**}(\ell_0) < \pi_0 < \pi^*(\ell_0)$ although such audit is not profitable in expected terms.

2.5 Conclusion

This article aimed at characterizing the learning dimension of auditing when there is a repeated interaction between auditor and auditee. The insurance claim fraud problem with potentially dishonest service providers was an application example, but the same question arises in many other settings, such as tax audits and, more generally, the verification of compliance with law. In our model, the insurer has imperfect information about the service providers' type, and, as in the machine learning multi-armed bandit approach, he extends his audit activity beyond the desire for immediate short-term gain. Compared to a myopic strategy only focusing on short term profit, the longsighted insurer faces an inter-temporal trade-off between the immediate gain from fraud detection, and the future profit made possible by more intense auditing. This learning effect leads the longsighted insurer to increase his monitoring efforts and to put some individually unprofitable claims under scrutiny. This result remains valid when the setting is extended to a more general framework with claims of varying size: the learning effect shifts the frontier in the belief-claim size space, beyond which an audit should be performed.

These results may be extended in many directions that would be worth exploring. Firstly, we have limited ourselves to a simple two-period model. Extending our analysis to an arbitrary number of periods would allow us to take into account the possibility to exclude and replace service providers when their dishonesty becomes very probable, and also to analyze the convergence features of our model when the number of periods is large. A parallel could also be drawn between our problem and the so-called greedy/ ϵ -greedy strategies in bandit problems, when the latter outperform the former in the long-run (see Chapter 2 in [Sutton and Barto \(1998\)](#)). Another interesting issue would consist in considering the case where service providers are concerned by multiple claims at each period. The auditing strategy would have to specify how many claims will be monitored. Audit costs may also reflect a potential source of heterogeneity between claims which could lead the insurer to abstain from auditing some claims with high audit costs. Furthermore, strategic defrauders could resort to manipulation of auditing costs as a protective device against auditing (see [Picard \(2000\)](#)). Most importantly, our model postulates an exogenous fraud rate reflected in the frequency of invalid claims for honest and dishonest service providers. Endogenizing

the frequency of invalid claims would be of particular interest, in order to study the strategic interaction between insurers, service providers and claimants, in a setting where learning and deterrence effects would coexist. Finally, the paper highlights the relevance of more intense investigations during the first stages of repeated interactions between principal and agent, and this conclusion is far more general than the insurance fraud problem. Analyzing whether this conclusion fits with actual monitoring processes would be an empirical extension that would be worth exploring, for better understanding the behavior of insurers facing claim fraud and in other contexts.

2.6 Appendix

Proof of Lemma 2.1. Let us write the expected profit from the auditing strategy $x(\cdot)$ at a single period

$$\Omega = \int_0^1 [\bar{p}(\pi) - c] x(\pi) f(\pi) d\pi.$$

The corresponding problem has a point-wise maximization structure, therefore:

$$\begin{aligned} x(\pi) &= 1 && \text{if } \pi > \pi^*, \\ x(\pi) &\in [0, 1] && \text{if } \pi = \pi^*, \\ x(\pi) &= 0 && \text{if } \pi < \pi^*, \end{aligned}$$

where

$$\pi^* = \frac{c - p_H}{p_D - p_H}.$$

□

Proof of Proposition 2.1. Let $\tilde{\pi}_1(\pi_0)$ be the period 1 belief as a function of prior belief π_0 . It is a random variable defined by

$$\tilde{\pi}_1(\pi_0) = \begin{cases} \pi_{1a} = A(\pi_0) & \text{with probability } q_a = x_0(\pi_0)\bar{p}(\pi_0), \\ \pi_{1b} = B(\pi_0) & \text{with probability } q_b = x_0(\pi_0)(1 - \bar{p}(\pi_0)), \\ \pi_{1c} = \pi_0 & \text{with probability } q_c = 1 - x_0(\pi_0), \end{cases}$$

with $\mathbb{E}_0[\tilde{\pi}_1(\pi_0)] = q_a\pi_{1a} + q_b\pi_{1b} + q_c\pi_{1c} = \pi_0$.

By a point-wise maximization argument, for all π_1 , the optimal auditing strategy $x_1^*(\pi_1)$ maximizes

$$\begin{aligned} &x_1[\bar{p}(\pi_1) - c] \\ \text{s.t. } &0 \leq x_1 \leq 1, \end{aligned}$$

and thus we have

$$x_1^*(\pi_{1i})[\bar{p}(\pi_{1i}) - c] \geq x_1^*(\pi_0)[\bar{p}(\pi_{1i}) - c] \text{ for } i = a, b, c.$$

Therefore

$$\begin{aligned}
 \mathbb{E}_0[\omega(\tilde{\pi}_1, x_1^*(\cdot)) | \pi_0] &= \sum_{i=a,b,c} q_i x_1^*(\pi_{1i}) [\bar{p}(\pi_{1i}) - c] \\
 &\geq x_1^*(\pi_0) \sum_{i=a,b,c} q_i [\bar{p}(\pi_{1i}) - c] \\
 &= x_1^*(\pi_0) \left[\bar{p} \left(\sum_{i=a,b,c} q_i \pi_{1i} \right) - c \right] \\
 &= x_1^*(\pi_0) [\bar{p}(\pi_0) - c] \\
 &= \omega(\pi_0, x_1^*(\pi_0)).
 \end{aligned}$$

This inequality is strict if $x_1^*(\pi_{1i}) \neq x_1^*(\pi_0)$ for some $i = a$ and/or $i = b$. □

Proof of Lemma 2.2. We have $\pi_1 \in \{B(\pi_0), \pi_0, A(\pi_0)\}$ depending on the period 0 scenario.

1. If $0 \leq \pi_0 < \pi_a$

(a) In all cases: $B(\pi_0) \leq \pi_0 \leq A(\pi_0) < A(\pi_a) = \pi^* \implies \pi_1 < \pi^*$.

2. If $\pi_a \leq \pi_0 < \pi^*$

(a) Invalid claim: $\pi_1 = A(\pi_0) \geq A(\pi_a) = \pi^* \implies \pi_1 \geq \pi^*$.

(b) Valid or No Audit

i. $\pi_1 = B(\pi_0) < \pi_0 < \pi^* \implies \pi_1 < \pi^*$.

ii. $\pi_1 = \pi_0 < \pi^* \implies \pi_1 < \pi^*$.

3. If $\pi^* \leq \pi_0 < \pi_b$

(a) Invalid or No Audit

i. $\pi_1 = A(\pi_0) > \pi_0 \geq \pi^* \implies \pi_1 \geq \pi^*$.

ii. $\pi_1 = \pi_0 \geq \pi^* \implies \pi_1 \geq \pi^*$.

(b) Valid claim:

i. $\pi_1 = B(\pi_0) < B(\pi_b) = \pi^* \implies \pi_1 < \pi^*$.

4. If $\pi_b \leq \pi_0 \leq 1$

(a) In all cases: $B(\pi_b) = \pi^* \leq B(\pi_0) \leq \pi_0 \leq A(\pi_0) \implies \pi_1 \geq \pi^*$.

The period 1 audit decision represented in Table 2.1 follows from the fact that there is an audit at period 1 if and only if $\pi_1 \geq \pi^*$. □

Proof of Lemma 2.3. Let $\Pi_0(\pi_0)$ be the expected inter-temporal net proceeds of auditing an SP of type π_0 . Let (the random variable) $\tilde{\pi}_1$ be the updated belief at period 1. We have

$$\begin{aligned}\Pi_0(\pi_0) &= x_0(\pi_0) \left[\bar{p}(\pi_0) - c \right] + \gamma \mathbb{E}_0 \left[x_1^*(\tilde{\pi}_1) \left(\bar{p}(\tilde{\pi}_1) - c \right) \mid \pi_0 \right] \\ &= x_0(\pi_0) \left[\bar{p}(\pi_0) - c \right] + \gamma \mathbb{E}_0 \left[\mathbb{1}_{\{\tilde{\pi}_1 \geq \pi^*\}} \left(\bar{p}(\tilde{\pi}_1) - c \right) \mid \pi_0 \right].\end{aligned}$$

If $\pi_0 \in [0, \pi_a)$, we always have $\tilde{\pi}_1 < \pi^*$ and

$$\Pi_0(\pi_0) = x_0(\pi_0) \left[\bar{p}(\pi_0) - c \right].$$

If $\pi_0 \in [\pi_a, \pi^*)$, then $\tilde{\pi}_1 \geq \pi^*$ if audit reveals an invalid claim (i.e., $\tilde{\pi}_1(\pi_0) = A(\pi_0)$), which happens with probability $\bar{p}(\pi_0)x_0(\pi_0)$. Thus, in that case we have

$$\begin{aligned}\Pi_0(\pi_0) &= x_0(\pi_0) \left[\bar{p}(\pi_0) - c \right] + \gamma \bar{p}(\pi_0)x_0(\pi_0) \left[\bar{p}(A(\pi_0)) - c \right] \\ &= x_0(\pi_0) \left[\bar{p}(\pi_0) - c \right] + \gamma \bar{p}(\pi_0)x_0(\pi_0) \left[p_D^2 \frac{\pi_0}{\bar{p}(\pi_0)} + p_H^2 \frac{1 - \pi_0}{\bar{p}(\pi_0)} - c \right] \\ &= x_0(\pi_0) \left[\bar{p}(\pi_0) - c \right] + \gamma x_0(\pi_0) \left[p_D^2 \pi_0 + p_H^2 (1 - \pi_0) - \bar{p}(\pi_0)c \right] \\ &= x_0(\pi_0) \left[\bar{p}(\pi_0) - c + \gamma \left[p_D^2 \pi_0 + p_H^2 (1 - \pi_0) - \bar{p}(\pi_0)c \right] \right].\end{aligned}$$

If $\pi_0 \in [\pi^*, \pi_b)$, then $\tilde{\pi}_1 \geq \pi^*$ if audit reveals an invalid claim (i.e., $\tilde{\pi}_1(\pi_0) = A(\pi_0)$ with probability $\bar{p}(\pi_0)x_0(\pi_0)$) or if there is no audit (i.e., $\tilde{\pi}_1(\pi_0) = \pi_0$ with probability $1 - x_0(\pi_0)$).

Hence

$$\begin{aligned}\Pi_0(\pi_0) &= x_0(\pi_0) \left[\bar{p}(\pi_0) - c \right] + \gamma \left[\bar{p}(\pi_0)x_0(\pi_0) \left[\bar{p}(A(\pi_0)) - c \right] + (1 - x_0(\pi_0)) \left[\bar{p}(\pi_0) - c \right] \right] \\ &= x_0(\pi_0) \left[\bar{p}(\pi_0) - c \right] + \gamma \left[\bar{p}(\pi_0) - c \right] + \gamma x_0(\pi_0) \left[p_D^2 \pi_0 + p_H^2 (1 - \pi_0) - \bar{p}(\pi_0)c - \bar{p}(\pi_0) + c \right] \\ &= \gamma \left[\bar{p}(\pi_0) - c \right] + x_0(\pi_0) \left(\bar{p}(\pi_0) - c + \gamma \left[p_D^2 \pi_0 + p_H^2 (1 - \pi_0) - \bar{p}(\pi_0)c - \bar{p}(\pi_0) + c \right] \right).\end{aligned}$$

If $\pi_0 \in [\pi_b, 1]$, then we always have $\tilde{\pi}_1 \geq \pi^*$, and thus

$$\Pi_0(\pi_0) = \gamma \left[\bar{p}(\pi_0) - c \right] + x_0(\pi_0) \left[\bar{p}(\pi_0) - c \right].$$

The expected net proceeds for an SP of type π_0 can therefore be written as:

$$\Pi_0(\pi_0) = C(\pi_0) + x_0(\pi_0)K(\pi_0),$$

where

$$K(\pi_0) = \bar{p}(\pi_0) - c + \gamma H(\pi_0),$$

and functions $C(\cdot)$ and $H(\cdot)$ are given in Table 2.3.

Table 2.3 – DEFINITION OF $C(\cdot)$ AND $H(\cdot)$.

π_0	$C(\pi_0)$	$H(\pi_0)$
$[0, \pi_a)$	0	0
$[\pi_a, \pi^*)$	0	$p_D^2 \pi_0 + p_H^2 (1 - \pi_0) - \bar{p}(\pi_0)c$
$[\pi^*, \pi_b)$	$\gamma [\bar{p}(\pi_0) - c]$	$p_D^2 \pi_0 + p_H^2 (1 - \pi_0) - \bar{p}(\pi_0)c - [\bar{p}(\pi_0) - c]$
$[\pi_b, 1]$	$\gamma [\bar{p}(\pi_0) - c]$	0

We obtain

$$\begin{aligned} \Omega_0(x_0(\cdot)) + \gamma \Omega_1^*(x_0(\cdot)) &= \int_0^1 \Pi_0(\pi_0) f_0(\pi_0) d\pi_0 \\ &= \int_0^1 [C(\pi_0) + K(\pi_0)x_0(\pi_0)] f_0(\pi_0) d\pi_0. \end{aligned}$$

The piecewise linearity of $K(\pi_0)$ comes from the fact that $\bar{p}(\pi_0)$ is linear in π_0 . In addition,

$$\begin{aligned} H(\pi_a) &= p_D^2 \pi_a + p_H^2 (1 - \pi_a) - \bar{p}(\pi_a)c \\ &= \bar{p}(\pi_a) [\bar{p}(A(\pi_a)) - c] \\ &= \bar{p}(\pi_a) [\bar{p}(\pi^*) - c] \\ &= 0. \end{aligned}$$

Thus

$$K(\pi_a) = \bar{p}(\pi_a) - c = \lim_{\pi \rightarrow \pi_a^-} K(\pi).$$

Notice also that

$$\begin{aligned} K(\pi^*) &= p_D^2 \pi^* + p_H^2 (1 - \pi^*) - \bar{p}(\pi^*)c - [\bar{p}(\pi^*) - c] \\ &= p_D^2 \pi^* + p_H^2 (1 - \pi^*) - \bar{p}(\pi^*)c \\ &= \lim_{\pi \rightarrow \pi^{*-}} K(\pi). \end{aligned}$$

Finally, by definition of $\pi^* = B(\pi_b)$, we have

$$\begin{aligned} c &= \bar{p}(B(\pi_b)) \\ &= p_D \frac{(1 - p_D)\pi_b}{1 - \bar{p}(\pi_b)} + p_H \left(1 - \frac{(1 - p_D)\pi_b}{1 - \bar{p}(\pi_b)}\right), \end{aligned}$$

and

$$p_D^2 \pi_b + p_H^2 (1 - \pi_b) - \bar{p}(\pi_b) c = (\bar{p}(\pi_b) - c).$$

This implies

$$K(\pi_b) = \bar{p}(\pi_b) - c = \lim_{\pi \rightarrow \pi_b^-} K(\pi),$$

which proves that $K(\cdot)$ is continuous. Finally, from the definition of $K(\cdot)$

$$\begin{aligned} \forall \pi_0 \in [0, 1] \quad K(\pi_0) &= \bar{p}(\pi_0) - c + \gamma H(\pi_0) \\ &\geq \bar{p}(\pi_0) - c. \end{aligned}$$

We also have

$$\forall \pi_0 > \pi^* \quad K(\pi_0) \geq \bar{p}(\pi_0) - c > 0,$$

and

$$\forall \pi_0 \leq \pi_a \quad K(\pi_0) = \bar{p}(\pi_0) - c < 0.$$

□

Proof of Proposition 2.2. Point-wise maximization yields

$$\begin{aligned} x_0^*(\pi_0) &= 0 && \text{if } K(\pi_0) < 0, \\ x_0^*(\pi_0) &\in [0, 1] && \text{if } K(\pi_0) = 0, \\ x_0^*(\pi_0) &= 1 && \text{if } K(\pi_0) > 0. \end{aligned}$$

Therefore, the threshold $\pi^{**} \in [\pi_a, \pi^*)$ satisfies $K(\pi^{**}) = 0$ and

$$\begin{aligned} K(\pi^{**}) &= \bar{p}(\pi^{**}) - c + \gamma H(\pi^{**}) \\ &= \left(1 + \gamma(p_D + p_H - c)\right)(p_D - p_H)\pi^{**} + (p_H - c)(1 + \gamma p_H). \end{aligned}$$

Using $K(\pi^{**}) = 0$ gives

$$\pi^{**} = \pi^* \frac{1 + \gamma p_H}{1 + \gamma p_H + \gamma(p_D - c)} < \pi^*.$$

□

Proof of Lemma 2.4. Let $\Delta = [0, 1] \times [0, L_1]$. The optimal strategy at period 1 defined by $x_1^*(\cdot, \cdot) : \Delta \rightarrow [0, 1]$ is such that

$$x^*(\pi_1, \ell_1) = \begin{cases} 1 & \text{if } (\pi_1, \ell_1) \in \Delta^*, \\ 0 & \text{otherwise.} \end{cases}$$

The associated period 1 objective function is

$$\Omega_1(x_0(\cdot, \cdot), x_1^*(\cdot, \cdot)) = \Omega_1^*(x_0(\cdot, \cdot)),$$

and thus the inter-temporal objective is written as a function of $x_0(\cdot, \cdot)$

$$\Omega_0(x_0(\cdot, \cdot)) + \gamma \mathbb{E}_0 \left[\Omega_1^*(x_0(\cdot, \cdot)) \right].$$

Since the random variable $\tilde{\ell}_i$ is independent of the type, we may write

$$\Omega_0(x_0(\cdot, \cdot)) = \int \int_{(\pi_0, \ell_0) \in \Delta} \left[\bar{p}(\pi_0)\ell_0 - c \right] x_0(\pi_0, \ell_0) f_0(\pi_0) g_0(\ell_0) d\ell_0 d\pi_0.$$

There is an audit at period 1 if $(\pi_1, \ell_1) \in \Delta^*$, and thus we have

$$\begin{aligned} \Omega_1^*(x_0(\cdot, \cdot)) &= \int \int_{(\pi_1, \ell_1) \in \Delta} \left[\bar{p}(\pi_1)\ell_1 - c \right] x^*(\pi_1, \ell_1) g_1(\ell_1) f_1(\pi_1 | x_0(\cdot, \cdot)) d\pi_1 d\ell_1 \\ &= \int_{\pi_1 \in [0, 1]} \int_{\ell^*(\pi_1)}^{L_1} \left[\bar{p}(\pi_1)\ell_1 - c \right] g_1(\ell_1) f_1(\pi_1 | x_0(\cdot, \cdot)) d\ell_1 d\pi_1 \\ &= \int_{\pi_1 \in [0, 1]} \Phi(\pi_1) f_1(\pi_1 | x_0(\cdot, \cdot)) d\pi_1, \end{aligned}$$

where $\ell^*(\pi_1) = c/\bar{p}(\pi_1)$ and

$$\Phi(\pi_1) = \int_{\ell^*(\pi_1)}^{L_1} [\bar{p}(\pi_1)\ell_1 - c] g_1(\ell_1) d\ell_1.$$

Note that $\Phi(\pi_1)$ is the expected net proceeds at period 1 of auditing an SP with belief π_1 . Therefore, by analogy with Section 2.3 and using a point-wise maximization argument, the inter-temporal expected net proceeds of an SP characterized by (π_0, ℓ_0) are written as

$$\begin{aligned} \Pi_0(\pi_0, \ell_0) = x_0(\pi_0, \ell_0) & \left[\bar{p}(\pi_0)\ell_0 - c \right] + \gamma \left[x_0(\pi_0, \ell_0)\bar{p}(\pi_0)\Phi(A(\pi_0))\dots \right. \\ & \dots + x_0(\pi_0, \ell_0)(1 - \bar{p}(\pi_0))\Phi(B(\pi_0))\dots \\ & \left. \dots + (1 - x_0(\pi_0))\Phi(\pi_0) \right]. \end{aligned}$$

Rearranging the terms in $\Pi_0(\pi_0, \ell_0)$ yields

$$\Pi_0(\pi_0, \ell_0) = C(\pi_0) + K(\pi_0, \ell_0)x_0(\pi_0, \ell_0),$$

where

$$\begin{aligned} K(\pi_0, \ell_0) &= \bar{p}(\pi_0)\ell_0 - c + \gamma H(\pi_0), \\ H(\pi_0) &= \bar{p}(\pi_0)\Phi(A(\pi_0)) + (1 - \bar{p}(\pi_0))\Phi(B(\pi_0)) - \Phi(\pi_0), \\ C(\pi_0) &= \gamma\Phi(\pi_0). \end{aligned}$$

Simple calculations give

$$\Phi'(\pi_1) = (p_D - p_H) \int_{\frac{c}{\bar{p}(\pi_1)}}^{L_1} \ell_1 g_1(\ell_1) d\ell_1 > 0,$$

and

$$\Phi''(\pi_1) = (p_D - p_H)^2 \frac{c^2}{\bar{p}(\pi_1)^3} g_1\left(\frac{c}{\bar{p}(\pi_1)}\right) > 0.$$

Thus, $\Phi(\cdot)$ is increasing and convex. Using $\bar{p}(\pi_0)A(\pi_0) + (1 - \bar{p}(\pi_0))B(\pi_0) = \pi_0$ gives $H(\pi_0) > 0$ if $\pi_0 \in (0, 1)$. In addition, $A(1) = B(1) = 1$ and $A(0) = B(0) = 0$ imply $H(0) = H(1) = 0$. \square

Proof of Proposition 2.3. Lemma 2.4 shows that $x_0^*(\pi_0, \ell_0) = 1$ if $K(\pi_0, \ell_0) > 0$ and that $x_0^*(\pi_0, \ell_0) = 0$ if $K(\pi_0, \ell_0) < 0$. Let $\pi_0 \in (\pi^*(\ell_0), 1)$. Since $H(\pi_0) > 0$ and $\bar{p}(\pi_0)\ell_0 - c \geq 0$, and

$H(1) = 0$ and $\bar{p}(1)\ell_0 - c > 0$, we have

$$K(\pi_0, \ell_0) > 0.$$

This implies that Δ^* is included in the optimal auditing set at period 0. In addition, $K(\pi^*(\ell_0), \ell_0) > 0$ implies, by continuity of $K(\cdot)$, that there exists $\pi^{**}(\ell_0)$ smaller than $\pi^*(\ell_0)$ such that

$$K(\pi_0, \ell_0) > 0 \quad \forall \pi_0 \in (\pi^{**}(\ell_0), \pi^*(\ell_0)).$$

We deduce that

$$x_0^*(\pi_0, \ell_0) = 1 \text{ if } \pi_0 > \pi^{**}(\ell_0).$$

□

Chapter 3

Should I Stalk or Should I Go? An Auditing Exploration/Exploitation Dilemma ¹

¹This chapter a joint work with Pierre Picard. We are grateful to conference participants at APRIA, ARIA and ESEM annual meetings, as well as seminar participants at CREST for helpful discussions and comments.

Abstract

We consider a fraud inspection problem where service providers are central to the fraud generating process, either as the main protagonists or as colluding third parties. Because interactions are repeated between the auditor (insurer, tax collector, environmental regulation agency, etc.) and auditees (doctors, tax preparers, waste management subcontractors, etc.), auditing behaves as a learning mechanism to separate the wheat (honest agents) from the chaff (defrauders). We analyze a Bayesian inspector's dynamic auditing problem in the face of fraud, and characterize its optimal strategy as a strategic exploration/one-armed bandit one. The insurer faces the well-known reinforcement learning exploration/exploitation trade-off between gathering information for higher future profits (exploration) and prioritizing immediate profits (exploitation). We then derive optimal auditing strategies with multiple auditees and capacity constraints as the solution to a k-armed bandit problem. We finally investigate the extents to which learning occurs under optimality in terms of how much information is obtained and how quickly it is obtained.

JEL Classification Numbers: D82, D83, G22, L21.

Keywords: Fraud, Dynamic optimal auditing, Information acquisition, Armed bandit.

3.1 Introduction

When an auditor can monitor an auditee's action at some cost, the auditing decision comes down to balancing between investigation-generated gains and losses. These gains and losses usually include the administrative costs of auditing, the proceeds from recovering undue payments in case of a fraud-revealing audit, and the profits drawn from a deterrence effect when defrauders are threatened to be spotted. But another non-immediate (and less scrutinized) motivation for monitoring may arise: information acquisition. When auditees are heterogeneous and the auditor's optimal action is not type-constant, the latter could benefit from learning more about the auditee. For example, in insurance, policyholders may differ in their propensity to defraud: it would be optimal for the insurer to audit "low morality" types more since they defraud frequently enough to amortize auditing costs, with significant deterrence effects. On the contrary, "high morality" types usually submit invalid claims by mistake, rendering them impervious to incentives, and unfrequently enough for net proceeds of auditing to be negative. While this morality is not directly observable, the repeated fraud patterns may convey information about the true type, provided an audit is conducted. This situation is obviously not restricted to insurance fraud and is relevant to multiple settings where fraud and auditing can take place. Tax evasion, environmental compliance and subcontractor fraud all fall in its scope and share a common feature, that is the repeated character (annually, monthly or even daily) of auditor-auditee interactions.² But an audit is costly, and gains from uncertainty reduction may not be profitable. When are they then?

To answer this question, we consider a dynamic programming problem where a decision maker (DM), the auditor, faces a non-strategic auditee who takes an action that may be invalid. We willingly restrict the occurrence of fraud to a simple non-strategic stochastic process in order to single out the information acquisition effect and exclude the deterrence one. The DM can choose to inspect the auditee at each period and reveal whether there was fraud. As a Bayesian, he acquires information by updating his belief about the auditee's type, conditionally on the outcome of the auditing decision. Our goal is to understand if this learning opportunity alters the DM's myopic optimal auditing policy and how. In particular, we are interested in the role of the time horizon (finite or infinite), the time preference and the efficiency of the information acquisition mechanism in reducing uncertainty.

Our main result states that this learning opportunity induces the DM to extend the inspection

²The frequency of interactions can also be explained by the existence of intermediaries who provide services to several agents: e.g., car repairers or health service providers in insurance, tax preparers for tax evasion

target set to some auditees whose immediate auditing proceeds are negative. In other words, the DM faces a trade-off between not auditing to avoid losing money at the expense of not learning anything, and auditing with a negative expected immediate proceed while acquiring knowledge for more accurate future audits. The DM's inter-temporal strategy takes the shape of a cut-off in the belief space, whereby he audits all agents above the threshold and ignores the ones below. With a finite-horizon, this cut-off is increasing in time, implying that more exploration takes place early on in the relationship: since the information acquired is about a time-invariant type, any information obtained at a given period is useful for all remaining periods, and, the further the horizon, the easier it is to amortize the one-time information acquisition cost. At the last period, the threshold is the one-period myopic one and excludes auditees with negative auditing proceeds. With an infinite horizon, the cut-off is unsurprisingly constant, lower than the myopic/no-learning threshold, and corresponds to the limit of the finite-horizon first period belief threshold when the total number of periods tends to infinity. As for the time preference, a more patient DM will be willing to explore more to enjoy the future information-generated additional profits. More importantly, when the DM is indifferent between the present and the future, i.e., when the discount rate tends to one, total exploration takes place in the infinite horizon setting: all beliefs are audited, as it takes a finite time and cost to reach a given level of belief precision, and an undiscounted infinite number of periods are left to take optimal actions and amortize the initial finite cost.

The intuitive concept underlying this set of results is the classical Exploration versus Exploitation dilemma: the DM must choose between playing his best option according to his current knowledge at the risk of being mistaken, or diversify his actions early on to refine his beliefs and take more accurate actions later on. This dilemma arises particularly in Reinforcement Learning problems, in particular Bandit problems. We establish that our dynamic programming problem is nothing but a Strategic Exploration/One-Armed Bandit one, i.e. a bandit with a risky arm (auditing) and a safe arm (not auditing). The value function is then analogous to the well-known Gittins index, which allows us to extend our main result to the case of multiple auditees and capacity constraints. Another interpretation of our problem, in line with the One-Armed Bandit approach, is that of an optimal stopping problem: the cut-off threshold defines a stopping time as the moment when the belief process crosses the threshold. In our fraud inspection problem, auditing has the peculiarity of being both the learning action and the profit generating action. Therefore, the DM's cut-off decision comes down to a trade-off between stopping early enough in case of a honest type (low stopping time/large threshold), and not excluding a dishonest type in

case of mistaken initial beliefs (large stopping time/low threshold). In contrast to our approach, the classical setting is one where the learning action (e.g., attention allocation, price revelation) is chosen first, then the profit generating action (buying, selling, etc.) is taken, and the DM's decision is the stopping time, which consequently never goes to infinity.

A second set of results elaborates on the specificity of the information acquisition process and its implications in terms of optimal auditing. We start by answering the question of when (in)complete/(in)adequate learning occurs, as defined by [Easley and Kiefer \(1988\)](#) and [Aghion et al. \(1991\)](#), and show that the answer lies in the distribution of optimal stopping times. Then, we focus on the efficiency of information gathering, i.e., the speed of convergence of the belief process to the true state. To do so, we first exploit the (sub/super)-martingale feature of the Bayesian belief process and focus on the Doob-Meyer predictable component.

These features relate to several strands of the academic literature. Our approach relies on the seminal work of [Townsend \(1979\)](#) and [Gale and Hellwig \(1985\)](#) on Costly State Verification where a principal can costly reveal some hidden information before taking the action. We extend the literature on optimal auditing³ by characterizing the scheme of optimal strategies under dynamic information acquisition in a multi-period setting.⁴ One important aspect through which our conclusions differ from previous ones is that our cutoff is in a belief space rather than on monetary amounts associated to fraud. Our main result bears a striking resemblance with that of [Dionne, Giuliano and Picard \(2008\)](#) as the auditor is induced to monitor auditees with negative expected payoff, although for different reasons. For us, this is motivated by the perspective of future profits due to sequential learning, while the motivation stems from a static deterrence effect in their case. This dynamic aspect with information gathering links our work to the literature on sequential experiments and optimal stopping, following the seminal papers by [Wald \(1945\)](#) and [Arrow, Blackwell and Girshick \(1949\)](#). These optimal stopping problems are intrinsically related to bandit ones⁵ and our model is easily reinterpreted as a One-Armed Bandit problem, bringing our work close to the literature on Bandits in economics ([Bergemann and Välimäki \(2008\)](#)) and on Reinforcement Learning ([Sutton and Barto \(1998\)](#)). More generally, it relates to the literature on dynamic information acquisition by a decision maker, as in [Easley and Kiefer \(1988\)](#) and [Aghion et al. \(1991\)](#).

The rest of the paper is organized as follows: Section 4.2 introduces the model. Section 3.3 characterizes the inter-temporal optimal auditing strategy and its interpretations in terms of

³see [Picard \(2013\)](#) for a comprehensive review.

⁴This work extends our previous analysis of the two-period case, see [Aboutajdine and Picard \(2018\)](#).

⁵see [Ferguson \(2004\)](#).

bandits and optimal stopping. Section 3.4 focuses on the extents to which learning by auditing happens and its informativeness as a function of the problem's parameters. Section 3.5 concludes. All proofs are in the Appendix.

3.2 The Model

3.2.1 Setting

We consider a world inhabited by two protagonists: an auditor and an auditee. The fundamental aspect of their interaction in this context is that it is repeated in time. Practically, it relates to the fact that many interactions occur between an auditing party and a service provider (SP). The SP can be the direct entity contractually related to the auditor, such as subcontractors who are major suppliers of a client firm or public institution. The SP can instead act as a third party to perform a commercial act and/or certify actions, such as health service providers (dentists, opticians, etc.) or tax preparers. These party SPs handle therefore the cases of several policyholders of the same insurer or of several taxpayers, thus the frequent interaction with the auditor.⁶ Hereafter, we restrict the auditee to be an SP, and use, for the sake of simplicity and without loss of generality, the terminology of insurance fraud. The auditor becomes the insurer, the auditee is an SP, and the SP's action is to submit a claim.

The SP channels one and only one claim at each period with value 1, each claim being either valid or invalid. The invalidity of a claim may result from honest mistakes or actual ill-disposed voluntary intent to defraud. The probability of submitting invalid claims depends on the SP's intrinsic type, with the SP being either honest (type H , transmits invalid claims only involuntarily) or dishonest (type D , transmits invalid claims both voluntarily and involuntarily). Because of this, a D type is more prone to submitting invalid claims than an H type, and we consider that he does so with probability p_D , while this probability is p_H for an H type, with $p_H < p_D$. Hence, the SP is a non-strategic⁷ entity and his actions are represented by an exogenous Bernoulli process of parameter p_i .

Types are time-invariant private information, unobservable to the insurer. But the latter, based on his experience and his observations, has a subjective belief π that the SP is of type D . In other words, the insurer attaches probability π to the fact that the SP is of type D . An SP

⁶We know for a fact that anti-fraud efforts in health insurance focus on service providers, as their collaboration is a necessary condition for fraud to occur, and they may even defraud without the insureds knowledge. The case of tax preparers is less obvious, but [Boning et al. \(2018\)](#) suggest that they can play an important role in anti-fraud efforts.

⁷Reminder: this is important to isolate the information effect from the deterrence effect.

to whom the auditor assigns a prior π will therefore, from the insurer's point of view, submit an invalid claim with probability $\bar{p}(\pi) = (1 - \pi)p_H + \pi p_D$. The insurer may audit a claim and reveal its true status, valid or invalid: when invalid, the auditor gets back the illegitimate amount of 1. This decision is represented by the choice of an auditing strategy $x \in [0, 1]$, which is the probability with which the auditor will investigate the claim submitted by the SP. Since the only information available to the auditor is the prior, the auditing strategy is a function of π .

It costs $c \in (p_H, p_D)$ to perform an audit, inducing a net proceed of $1 - c$ when a claim is found invalid, and $-c$ when it is found valid. In expectation, the net proceed of auditing an SP with prior π is $\bar{p}(\pi) - c$. In particular, if the type was known, auditing a dishonest SP would yield an expected net proceed of $p_D - c > 0$ while auditing an honest one would yield $p_H - c < 0$. This difference in the profitability of auditing between types of SPs is what may motivate the insurer to acquire information.

Finally, interactions take place over an arbitrary number of periods $T \geq 1$.⁸ Each period is indexed by $t \in \{1, 2, \dots, T\}$ and the initial period $t = 1$ is the beginning of the insurer-SP relationship. The period 1 initial prior may be considered as initialized at some arbitrary value (e.g. 0.5 if no relevant information is available) or based on some other observable characteristics of the SP. The total number of periods over which the interactions take place may be interpreted as the "lifespan" of an SP. For example, if an insurer checks the channeled claims every month and a pharmacist works for 30 years (i.e., 360 months), then $T = 360$. Whenever necessary, the time-dependent variables of interest will be indexed by both the period at which they are considered and the total number of periods, i.e., by $(t, T) \in \{1, \dots, T\} \times \mathbb{N}^*$. Otherwise, we will restrain the indexation to $t \in \{1, \dots, T\}$.

3.2.2 Information Acquisition

Period t audits also allow the insurer to update his beliefs at the beginning of period $t + 1$. Depending on whether an audit has been performed and, if so, whether the claim was valid or invalid (*Val* and *Inv*, respectively), posterior beliefs $\tilde{\pi}_{t+1}$ are deduced from initial beliefs π_t through Bayes' Law:

$$\tilde{\pi}_{t+1} = \begin{cases} \mathbb{P}(D|audit, Inv) = \varphi_{AI}(\pi_t) = \frac{p_D \pi_t}{\bar{p}(\pi_t)}, \\ \mathbb{P}(D|audit, Val) = \varphi_{AV}(\pi_t) = \frac{(1-p_D)\pi_t}{1-\bar{p}(\pi_t)}, \\ \mathbb{P}(D|no\ audit) = \varphi_{NA}(\pi_t) = \pi_t, \end{cases} \quad (3.1)$$

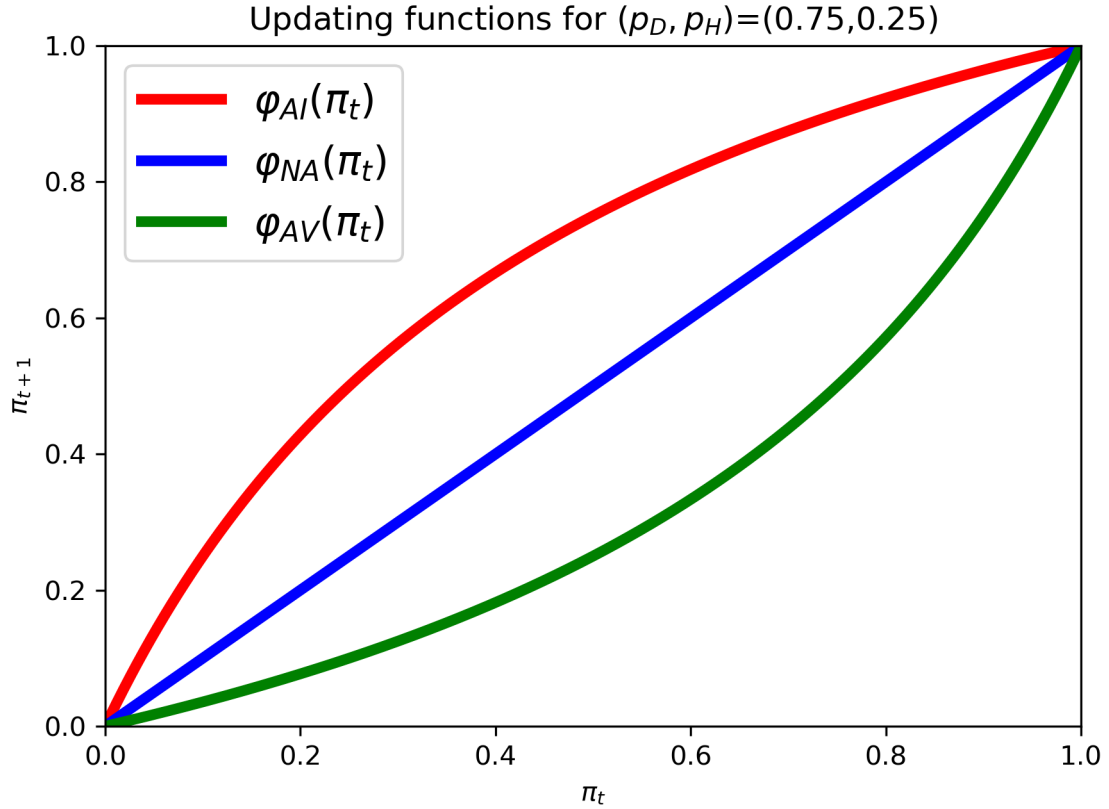
⁸The model is now a generalization of the two-period model considered by [Aboutajdine and Picard \(2018\)](#).

with

$$\begin{aligned} \varphi_{AV}(\pi_t) &< \pi_t < \varphi_{AI}(\pi_t), \\ \varphi'_{AI} &> 0, \quad \varphi''_{AI} < 0, \\ \varphi'_{AV} &> 0, \quad \varphi''_{AV} > 0. \end{aligned} \tag{3.2}$$

The updating mechanism is such that an invalid claim increases the belief that an SP is dishonest

Figure 3.1 – UPDATING FUNCTIONS



($\pi_{t+1} = \varphi_{AI}(\pi_t) > \pi_t$) while a valid claim decreases it ($\pi_{t+1} = \varphi_{AV}(\pi_t) < \pi_t$). Either way, both cases imply that the insurer acquires information about the auditee's true type. On the contrary, not auditing leaves the beliefs unchanged since no relevant information is obtained.

As the setting is dynamic, the evolution of beliefs in times constitutes a stochastic belief process $\langle \pi_t \rangle$ with a natural filtration \mathcal{F}_t , and, because it is governed by Bayes Law, the process behaves as a martingale.

Lemma 3.1. *The belief process $\langle \pi_t \rangle$ is a martingale under the auditor's probability measure \mathbb{P}*

$$\mathbb{E}_{\mathbb{P}}[\tilde{\pi}_{t+1} | \mathcal{F}_t] = \pi_t.$$

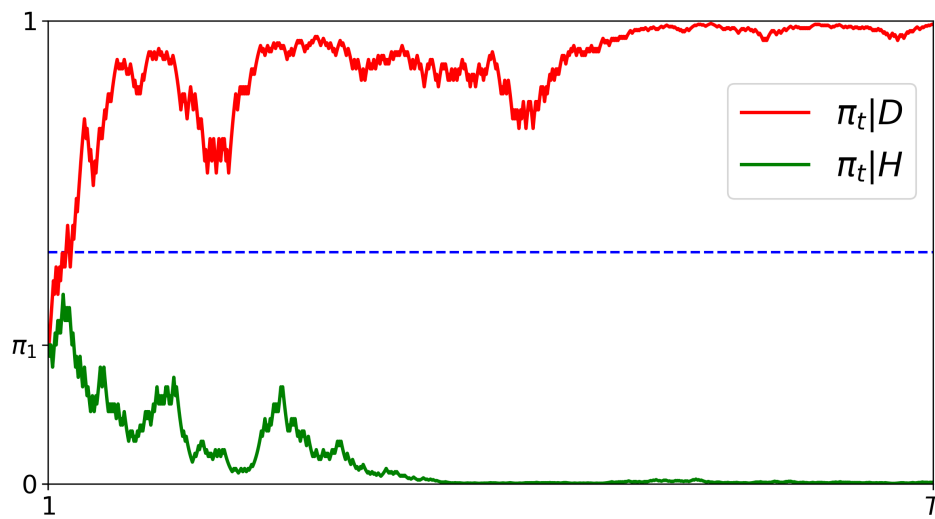
In addition, with some auditing (i.e. $x > 0$), under state of the world H (respectively D) and

the associated probability measures \mathbb{P}_H (resp. \mathbb{P}_D), the belief process $\langle \pi_t \rangle$ is a supermartingale (resp. a submartingale). With an infinite horizon, the process converges to the true belief about the state of the world

$$i - \mathbb{E}_H[\tilde{\pi}_{t+1}|\mathcal{F}_t] \leq \pi_t \text{ and } \pi_t \xrightarrow{\mathbb{P}_H} 0.$$

$$ii - \mathbb{E}_D[\tilde{\pi}_{t+1}|\mathcal{F}_t] \geq \pi_t \text{ and } \pi_t \xrightarrow{\mathbb{P}_D} 1.$$

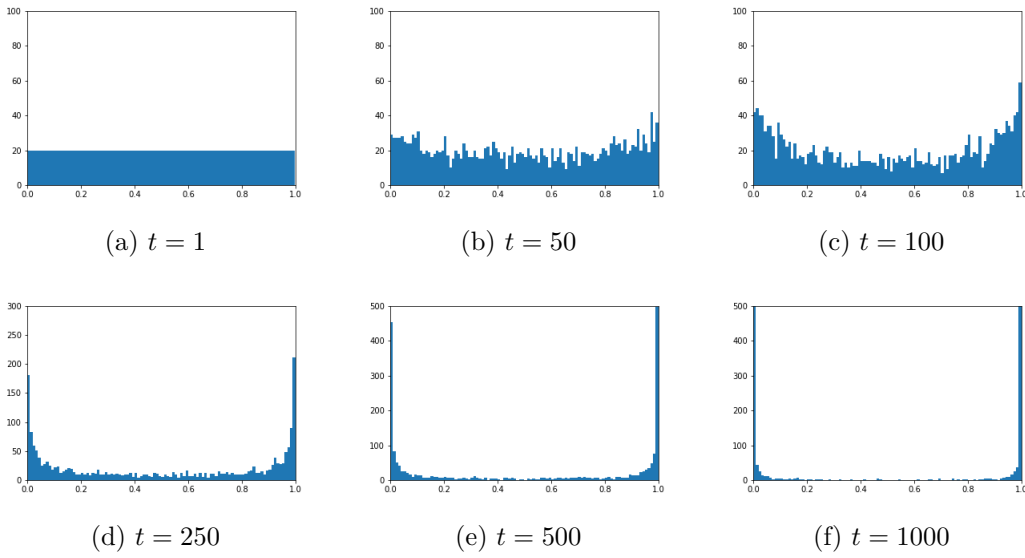
Figure 3.2 – SIMULATED CONVERGENCE OF PRIORS TO THE TRUE BELIEFS



Corollary 3.1. *The posterior $\tilde{\pi}_{t+1}$ is a mean-preserving spread of the prior π_t .*

Figure 3.2 illustrates the convergence of beliefs for each type, with the same initial belief π_1 . This convergence of the Bayesian beliefs to the truth under systematic auditing means that auditing acts as a separating tool to isolate dishonest SPs from honest ones. Figure 3.3 simulates over $T = 1000$ periods the trajectories of beliefs for a population of 2000 SPs, half honest half dishonest. At period 1, the population is uniformly distributed across priors, and a Bernoulli process is drawn for each SP's fraud behavior, resulting in a corresponding belief process realization under systematic auditing (i.e., $x = 1$). The progressive separation of types, as a convergence of beliefs to both ends of the spectrum, is such that, after enough time, dishonest SPs end up on the right, while honest ones are on the left.

In the following, we go back to a situation where the insurer faces one single SP.

Figure 3.3 – DISTRIBUTION OF PRIORS AFTER REPEATED AUDITS ($T = 1000$).


3.2.3 A Dynamic Programming Problem

3.2.3.1 The one-period/myopic problem

When $T = 1$, the problem is myopic, the auditor has no future profits to take into account, and information acquisition is not relevant. The auditor has objective function

$$\Omega_T(\pi, x) = u(\pi, x) = (\bar{p}(\pi) - c)x,$$

and solves for the value function

$$V_T(\pi) = \max_{x(\pi) \in [0,1]} \Omega_T(\pi, x).$$

Because of the linearity of the objective function in x , optimal auditing has a bang-bang solution characterized by a threshold π^+ above which auditing proceeds are positive. In other words, the auditors targets only individually profitable claims. The value function is consequently piece-wise linear in π .

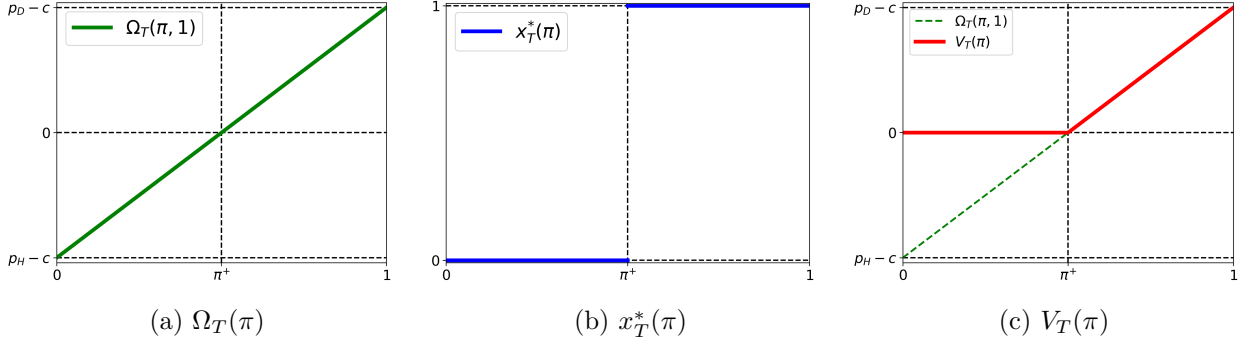
Lemma 3.2. *For $T = 1$, the optimal auditing strategy is bang-bang*

$$x_T^*(\pi) \begin{cases} = 0 & \text{if } \pi < \pi^+, \\ \in (0, 1) & \text{if } \pi = \pi^+, \\ = 1 & \text{if } \pi > \pi^+. \end{cases} \quad \text{where } \pi^+ = \frac{c - p_H}{p_D - p_H} \in (0, 1),$$

Figure 3.4 illustrates the myopic problem's objective function, optimal policy and value

function.

Figure 3.4 – MYOPIC OBJECTIVE, POLICY AND VALUE FUNCTIONS.



3.2.3.2 The multi-period problem

With a multi-period setting, the existence of future periods and their interactions with the present through the learning process alters the objective function at any $t < T$,

$$\Omega_t(\pi_t, x_t) = u(\pi_t, x_t) + \delta \mathbb{E}_t[V_{t+1}(\tilde{\pi}_{t+1})|x_t], \quad (3.3)$$

and

$$V_t(\pi_t) = \max_{x_t \in [0,1]} \Omega_t(\pi_t, x_t),$$

where δ represents time-discounting.⁹ Denote also with the superscript d the average discounted counterparts $\Omega_t^d = \frac{1}{\sum_{j=t}^T \delta^{j-t}} \Omega_t$ and $V_t^d = \frac{1}{\sum_{j=t}^T \delta^{j-t}} V_t$. Our problem is now a dynamic programming one with parameters $[\delta, u, \mathbb{P}(\tilde{\pi}_{t+1}), x]$.

In a finite horizon setting, this problem can be solved by backward induction. Therefore, the total number of periods T is important and we will occasionally index functions by $\{t, T\}$ rather than just $\{t\}$. The (piece-wise) linearity features of the objective/value functions persist through time and the corresponding optimal auditing strategies remain bang-bang. Lemma 3.3 characterizes all these features.

Lemma 3.3. *The objective function $\Omega_{t,T}(\pi, x)$ is*

i - linear in x , continuous and piece-wise linear in π ,

ii - $\frac{\partial \Omega_{t,T}}{\partial x} \Big|_{\pi}$ is continuous and increasing in π , $\frac{\partial \Omega_{t,T}}{\partial x} \Big|_{\pi=1} > 0$ and $\frac{\partial \Omega_{t,T}}{\partial x} \Big|_{\pi=0} < 0$,

⁹Our model may be extended to the case where the horizon T is random, and δ includes a termination probability.

iii - and there exists a threshold $\pi_{t,T}^* \in (0, 1)$ defined uniquely by $\left. \frac{\partial \Omega_{t,T}}{\partial x} \right|_{\pi_{t,T}^*} = 0$.

The optimal auditing strategy $x_{t,T}^*(\cdot)$ is bang-bang: $x_{t,T}^*(\pi) = \mathbb{1}_{\{\pi \geq \pi_{t,T}^*\}}$.

The value function $V_{t,T}(\pi)$ verifies

$$i - V_{t,T}(\pi) = \Omega_{t,T}(\pi, x_{t,T}^*(\pi)) = \max(0, \Omega_{t,T}(\pi, 1)),$$

ii - it is continuous and piece-wise linear in π ,

iii - it is convex in π (as the maximum of linear functions).

Figure 3.5 illustrates the Lemma 3.3 statements. Figure 3.5a illustrates the fact that optimal auditing is bang-bang. Figure 3.5b represents a possible shape of the intertemporal optimal auditing thresholds, with one threshold for each period. At this stage, we can't say anything more about these thresholds other than that they are lower than π^+ . Figure 3.5c illustrates the fact that the objective function at time t where the insurer audits with probability one is null for the corresponding threshold π_t^* . Finally, Figures 3.5d and 3.5e show how the value function is obtained either as the maximum between 0 and the objective function under certain auditing, or as the maximum of linear functions. At this stage of the paper, these linear functions appear in the proof of Lemma 3.3, but we don't know more about their signification.

At any time $t \in \{1, \dots, T-1\}$, the objective function in equation (3.3) differs from the myopic one through the second expectation term, which represents the impact of current auditing decision x_t on expected future auditing proceeds through its effects on beliefs. Lemma 3.4 describes this component's dependence on x_t .

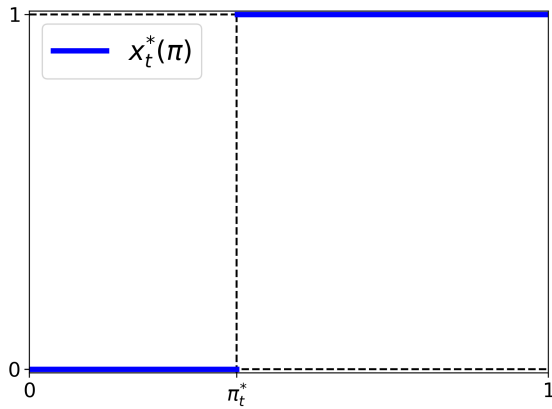
Lemma 3.4. *The component $\mathbb{E}_t[V_{t+1}(\tilde{\pi}_{t+1})|x_t]$ is increasing in x_t and always positive.*

In other words, current auditing increases expected future proceeds because of an exploration effect: the information gathered at the current period allows more accurate audits in the future. The objective function can therefore be decomposed into an exploitation and an exploration component

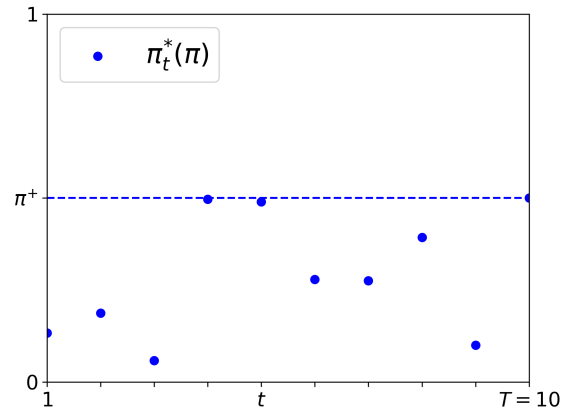
$$\Omega_t(\pi_t, x_t) = \overbrace{u(\pi_t, x_t)}^{\text{Exploitation}} + \delta \overbrace{\mathbb{E}_t[V_{t+1}(\tilde{\pi}_{t+1})|x_t]}^{\text{Exploration}}.$$

The myopic threshold $\pi^+ = \pi_{T,T}^*$ divides the belief space into two segments where exploitation and exploration effects may go in different directions (see Table 3.1): above π^+ , both components are increasing in x_t , while below π^+ , the exploitation effect is decreasing in x_t , discouraging auditing, while the exploration one is increasing in x_t , encouraging it.

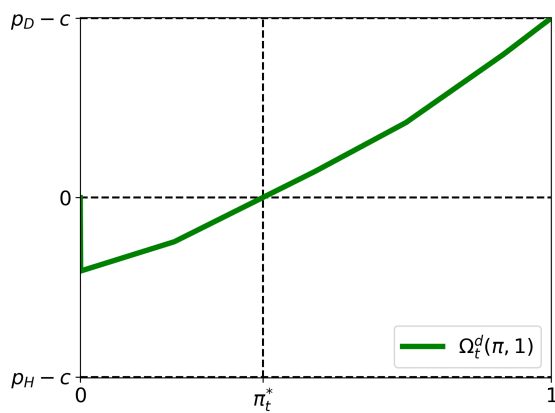
Figure 3.5 – TIME t POLICY, OBJECTIVE AND VALUE FUNCTIONS



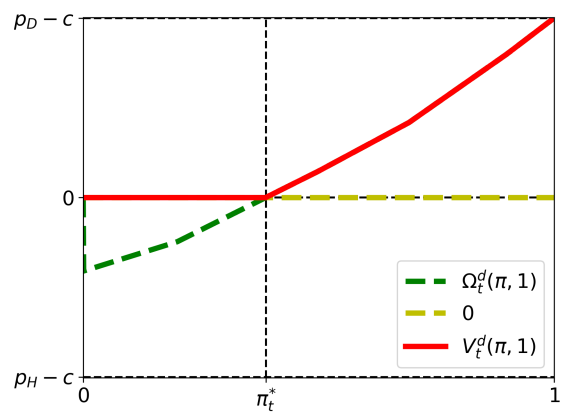
(a) OPTIMAL POLICY FUNCTION



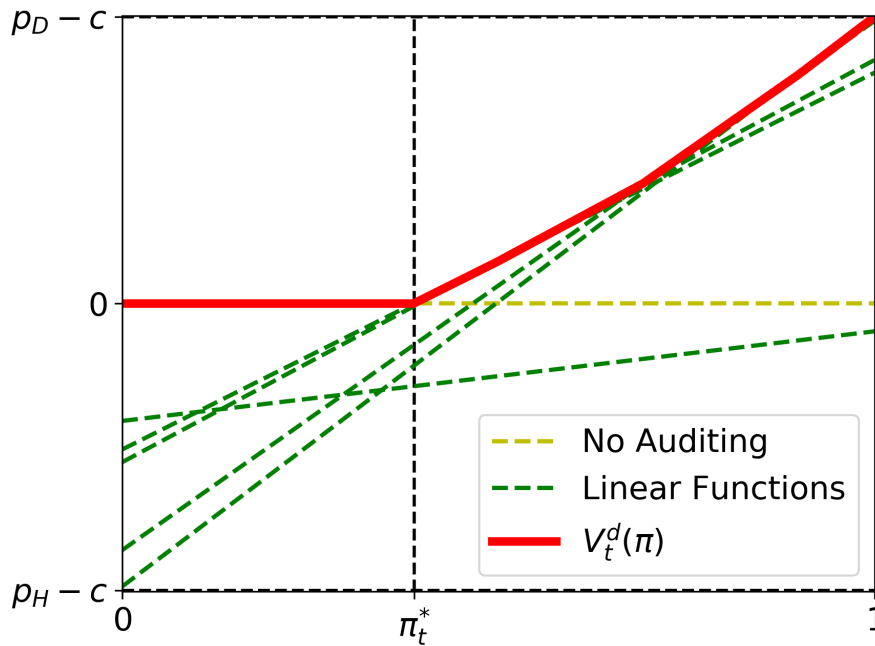
(b) AUDITING THRESHOLDS



(c) OBJECTIVE FUNCTION



(d) $V_t = \max(0, \Omega_t(\pi, 1))$



(e) $V_t = \max(\text{Linear functions})$

Table 3.1 – MARGINAL EFFECTS OF AUDITING

		$(0, \pi^+)$	$(\pi^+, 1)$
Exploitation	$\frac{\partial}{\partial x} u$	-	+
Exploration	$\frac{\partial}{\partial x} \mathbb{E}_t[V_{t+1}]$	+	+
Objective	$\frac{\partial}{\partial x} \Omega_{t,T}(\pi)$?	+

The exploration and exploitation effects are thus conflicting on $(0, \pi^+)$. The exploitation component is maximized for $\pi_t^* = \pi^+$. The exploration component is maximized for $\pi_t^* = 0$. We can therefore identify the sequence (π_t^*) , illustrated in Figure 3.5b, as a measure of the dynamic balance between exploration and exploitation. As illustrated in Figure 3.6, the closer π_t^* to 0, the more the auditor explores, and the closer π_t^* to π^+ , the more he exploits.

Figure 3.6 – LEVELS OF EXPLORATION/EXPLOITATION



Our goal hereafter is to understand how this optimal balance between exploration and exploitation behaves through time, i.e., to characterize the sequence (π_t^*) .

3.3 Exploration vs Exploitation: Auditing to Separate the Wheat from the Chaff

In this section, we present our main results about the evolution of optimal auditing thresholds as functions of time and discounting, and we relate our problem to classical learning problems. In Section 3.3.1, we answer our question of interest for the finite horizon case, then take the limit when T tends to infinity to characterize the infinite horizon solution in Section 3.3.2. Finally we interpret our problem as an optimal stopping one in Section 3.3.3 and as a Bandit one in 3.3.4.

3.3.1 Finite horizon optimal auditing

The inclusion of the exploration component at early periods encourages more auditing by decreasing the optimal thresholds. In addition, the larger the discount rate δ , the stronger the exploration effect. These results are stated in Proposition 3.1.

Proposition 3.1 (Finite Horizon Optimal Thresholds). *The existence of the learning opportunity changes the auditing efforts as*

i - the optimal thresholds sequence $(\pi_{t,T}^)_{t \in \{1, \dots, T\}}$ is strictly increasing, i.e., there is more exploration in the early stages of the relationship,*

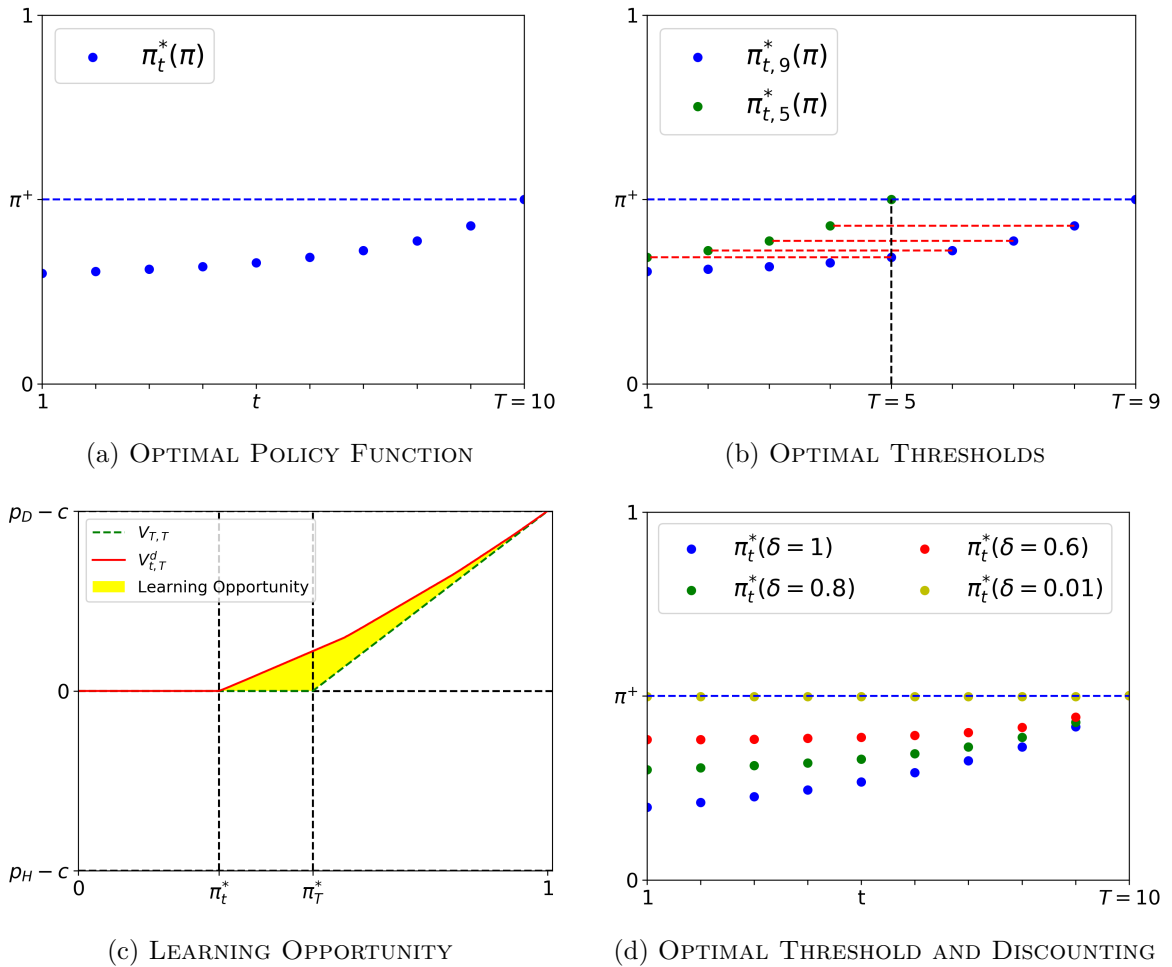
ii - the optimal level of exploration depends on the remaining number of periods $(\pi_{t,T}^ = f(T-t))$ and for all $i \in \{0, 1, \dots, t-1\}$, $\pi_{t,T}^* = \pi_{t-i, T-i}^* = \pi_{1, T-t+1}^*$,*

iii - the sequence of first period thresholds $(\pi_{1,T}^)_{T \in \mathbb{N}^*}$ is strictly decreasing,*

iv - the optimal levels of exploration are increasing with the DM's patience, i.e., for any $t \in \{1, \dots, T-1\}$:

$$\frac{\partial \pi_{t,T}^*}{\partial \delta} < 0 \quad \text{and} \quad \lim_{\delta \rightarrow 0} \pi_{t,T}^*(\delta) = \pi^+.$$

Figure 3.7 – FINITE HORIZON THRESHOLDS AND VALUE FUNCTION



Statement **i** means that, for a given number of periods, the earlier the audit, the larger the auditing set $(\pi_{t,T}^*, 1]$. Figure 3.7a illustrates this through the fact that the optimal thresholds are increasing in time. This is due to the fact that interacting for a longer time with auditees induces the auditor to learn more about them, even when their claims are believed to be non-profitable to audit at that point in time. Because information gained from the current period audit increases all subsequent periods' expected auditing proceeds, the generated gains cover the one-time preliminary expenses. In terms of exploration/exploitation, the auditor explores more at the beginning, hoping to identify dishonest types soon enough to exploit them later. Statement **ii** formalizes this dependence on the remaining number of periods. As shown in Figure 3.7b, starting at period 1 with a horizon $T = 5$ yields the same optimal thresholds as starting at period 5 with a horizon $T = 9$: in both cases, there are 4 periods remaining, and $(\pi_{t,5}^*)_{t \in \{1, \dots, 5\}} = (\pi_{t,9}^*)_{t \in \{5, \dots, 9\}}$. As a consequence, the average discounted value function $V_{t,T}^d$ is larger than the myopic value function $V_{T,T}$, and the difference represents the learning opportunity, as illustrated in Figure 3.7c. Statement **iii** is a direct consequence of statements **i** and **ii** because of the backward induction nature of the solution, and conveys the idea that, at the beginning of the relationship, more exploration takes place if there are more periods remaining (i.e. T increases). Statement **iv** tells us that the more patient the insurer, the more exploration takes place as more weight is put on the remaining periods. However, if the auditor is not patient, optimal auditing sequences tend to the myopic/one-period threshold π^+ . Figure 3.7d illustrates this last statement.

Observe that the increasing optimal thresholds mean that there is more auditing at the beginning of the relationship. [Mittone \(2006\)](#) notes in a tax evasion experiment that auditing may act as an “educating” mechanism when applied early. Participants audited at the beginning show less propensity to defraud in subsequent periods, even if they are no longer audited. On the contrary, subjects not audited before the second part of the experiments continue defrauding and are hardly deterred. Our result provides an alternative information based motivation for intense auditing earlier in an auditor-auditee relationship.

3.3.2 Infinite horizon optimal auditing

With an infinite horizon, the optimal auditing strategy is characterized in Proposition 3.2 below.

Proposition 3.2 (Infinite Horizon Optimal Thresholds). *The infinite-horizon problem is obtained by taking the limit $T \rightarrow \infty$*

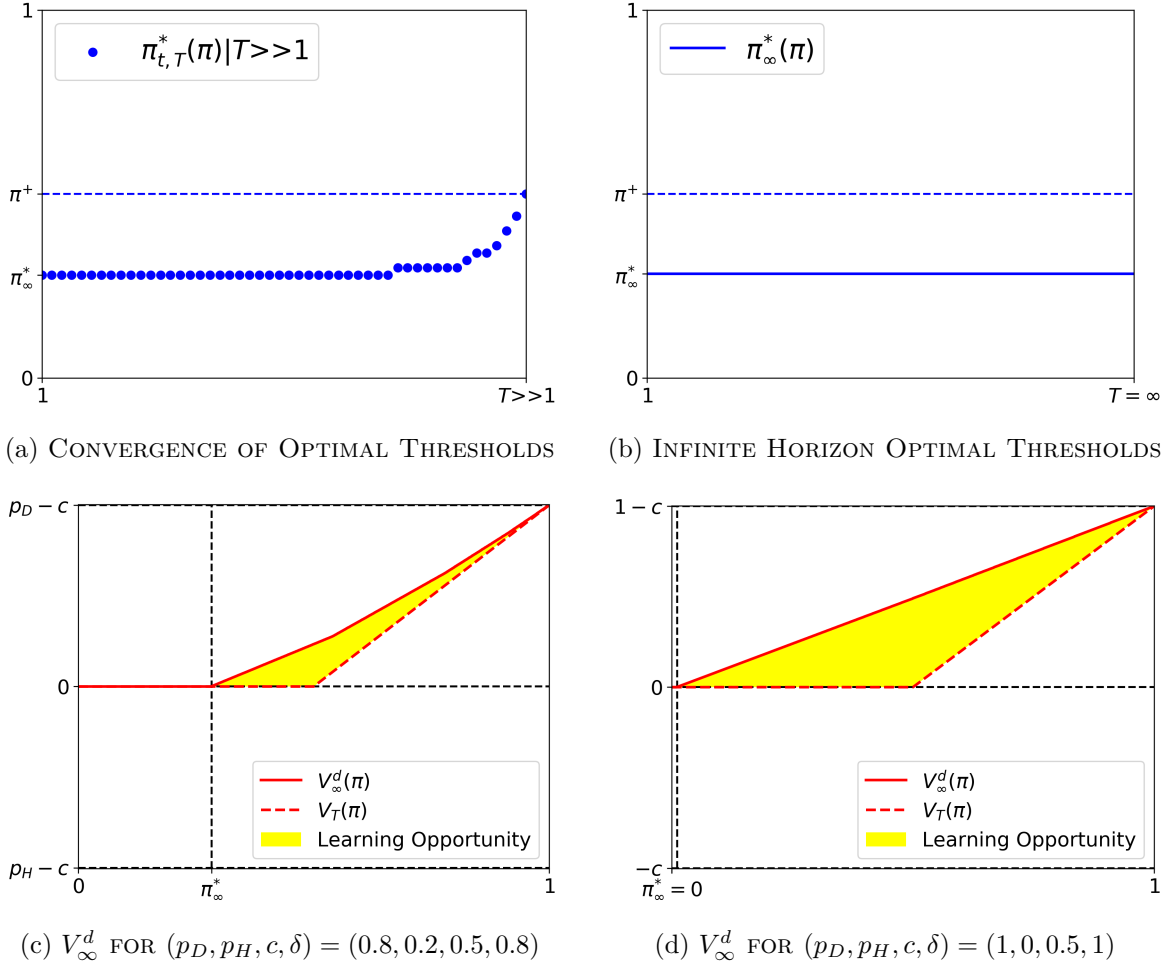
i - The optimal thresholds sequence $(\pi_{t,T}^*)_{t \in \mathbb{N}^*}$ is time independent, i.e. constant and equal to $\pi_\infty^* = \lim_{T \rightarrow \infty} \pi_{1,T}^* \geq 0$.

ii - The infinite horizon objective function $V_\infty(\pi)$ is convex in π as the supremum of an infinite countable family of linear functions.

iii - Total exploration occurs when the DM is patient enough

$$\lim_{\delta \rightarrow 1} \pi_\infty^*(\delta) = 0.$$

Figure 3.8 – INFINITE HORIZON THRESHOLDS AND VALUE FUNCTIONS



Statements *i* is illustrated in Figure 3.8a and Figure 3.8b, while statement *ii* is illustrated in Figure 3.8c. According to statement *iii*, a necessary and sufficient condition for all priors to be audited at the initial period when T becomes arbitrarily large is for the auditor to be patient enough, i.e., δ close enough to 1. The underlying intuition is that, while additional auditing allows to have refined information and better targeted audits, the lower the initial period prior,

the longer the insurer has to wait before he can derive positive proceeds for the corresponding SPs.

Example 3.1. Assume $(p_D, p_H) = (1, 0)$. In this context, types are revealed after only one period as $\varphi_{AI} = 1$ and $\varphi_{AV} = 0$. Then, the objective function is given by

$$\begin{aligned}\Omega_\infty(\pi, x) &= \left(\pi - c + \pi \sum_{i \geq 1} \delta^i (1 - c) \right) x \\ &= \left(\pi - c + \pi(1 - c) \frac{\delta}{1 - \delta} \right) x \\ &= \left(\pi \left(\frac{1 - \delta c}{1 - \delta} \right) - c \right) x.\end{aligned}$$

Then π_∞^* verifies $\pi \left(\frac{1 - \delta c}{1 - \delta} \right) - c = 0 \Rightarrow \pi_\infty^* = \frac{(1 - \delta)c}{1 - \delta c}$ and

$$V_\infty(\pi) = \left(\pi \left(\frac{1 - \delta c}{1 - \delta} \right) - c \right) \mathbb{1}_{\{\pi > \pi_\infty^*\}}.$$

Figure 3.8d illustrates the average discounted value function V_∞^d as a linear function when $\delta \rightarrow 1$, as linear combination with weights π and $1 - \pi$ between per period payoffs $1 - c$ and 0. Indeed, in this case, with probability π , the auditee is dishonest and the average future payoff per period is $1 - c$, and with probability $1 - \pi$ the average future payoff per period is 0, as auditing is stopped. Current period proceeds do not appear as they are amortized over an infinite undiscounted number of periods.

3.3.3 Dynamic auditing as an optimal stopping problem

A consequence of 3.1 and 3.2 is that once it is optimal to not audit at a given period, not auditing is optimal at all subsequent periods. This is due to the fact that our optimal auditing problem can be formulated as an optimal stopping problem, as shown in Proposition 3.3 below.

Proposition 3.3. Let τ be a stopping time with respect to $(\mathcal{F}_t)_{t \in \{1, \dots, T\}}$, where $T \in \mathbb{N} \cup \{\infty\}$, and define $W_{1, \tau}(\pi_1, \delta)$ as the expected proceeds of systematically auditing for τ periods, starting from a prior π_1 and with time discounting δ ,

$$W_{1, \tau}(\pi_1, \delta) = \mathbb{E} \left[\sum_{j=1}^{\tau-1} \delta^{j-1} u(\tilde{\pi}_j, 1) \right], \quad \text{and} \quad W_{1, 1}(\pi_1, \delta) = 0.$$

Then, for $\nu = \mathbb{1}\{\cup_{j \in \{1, \dots, T\}} \{\tilde{\pi}_j < \pi_j^*\}\}$, the optimal stopping time τ^* , defined as

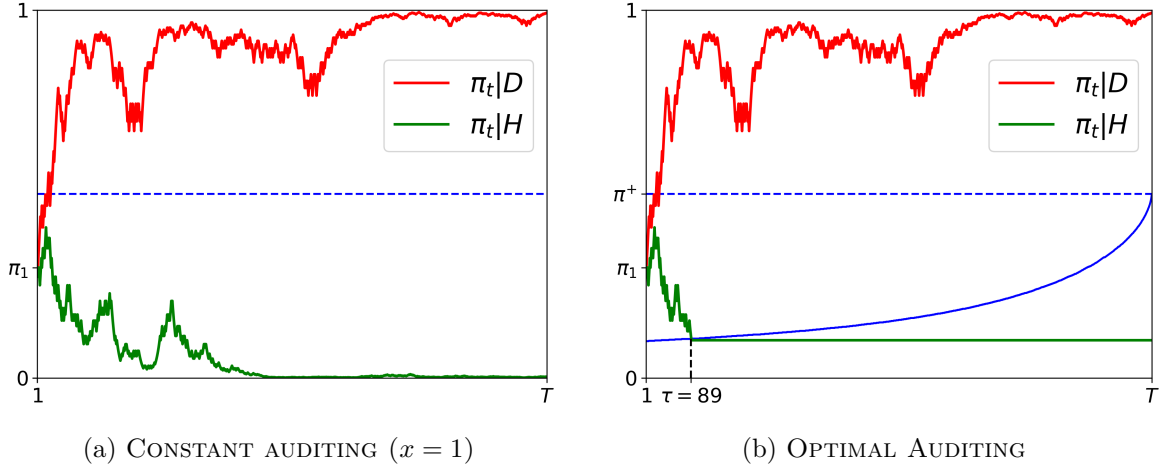
$$\tau^* = \nu \min \left\{ t \in \{1, \dots, T\} \mid \mathbb{1}\{\tilde{\pi}_t < \pi_t^*\} \right\} + (1 - \nu)(T + 1),$$

verifies

$$V_{1,T}(\pi_1, \delta) = W_{1,\tau^*}(\pi_1, \delta) = \max_{\tau} W_{1,\tau}(\pi_1, \delta).$$

Figure 3.9 illustrates the optimal stopping rule. When auditing is systematic, in Figure 3.9a, beliefs for both types converge towards the true belief. Under the optimal strategy, in Figure 3.9b, type H stops being audited after 89 periods, because the corresponding belief drops below the optimal threshold. Type D , on the contrary, continues being audited until the end. This optimal stopping rule defines an optimal stopping region in the belief space at each period t as $SR_t = \{\pi \mid \pi < \pi_t^*\}$. As soon as the belief process $\langle \pi_t \rangle$ generated under the optimal auditing strategy is lower than π_t^* , auditing and learning stop, and consequently the belief remains constant.

Figure 3.9 – OPTIMAL AUDITING AS AN OPTIMAL STOPPING RULE



This formulation as an optimal stopping problem allows us to better understand what are the linear functions of which $V_{1,T}$ is the maximum (Figure 3.5e). proposition 3.4 formalizes this idea.

Proposition 3.4. *There exists an integer $M \in \mathbb{N} \cup \{\infty\}$ and a partition $\cup_{m \in \{1, \dots, M\}} \Psi_m$ of $(0, 1)$ defined as*

$$(\pi, \pi') \in \Phi_m^2 \iff \forall \omega = (\omega_1, \dots, \omega_T), \tau^*(\omega, \pi) = \tau^*(\omega, \pi'),$$

where ω is a trajectory of the fraud generating process. Then, $V_{1,T}(\pi)$ is linear on each Ψ_m , and the family of linear functions of which $V_{1,T}$ is the maximum corresponds to the linear functions on Ψ_m .

Proposition 3.4 means that the value function is linear on intervals of beliefs for which auditing stops at the same time for all priors of the same interval. In other words, the trajectories of the belief process for a realization ω are such that the time t belief is either higher or lower than π_t^* , for all the initial beliefs belonging to the same interval.

In particular, with a finite horizon, one such interval is defined as $\{\pi | \pi > \varphi_{AV}^{t-T}(\pi_T^*)\}$. This is the set of beliefs that are so large that they are always audited, because updating will never bring them below an optimal threshold. In this case, the corresponding optimal stopping time is always equal to $T + 1$, and

$$\begin{aligned} V_{1,T}((\pi_1, \delta)) &= W_{1,T+1}(\pi_1, \delta) = \sum_{j=1}^T \delta^{j-1} \mathbb{E}[u(\tilde{\pi}_j, 1)] \\ &= \sum_{j=1}^T \delta^{j-1} u(\mathbb{E}[\tilde{\pi}_j], 1) \\ &= \sum_{j=1}^T \delta^{j-1} u(\pi_1, 1), \quad \text{from the martingale property.} \end{aligned}$$

This shows that for the highest beliefs, the average discounted value function is equal to the myopic value function.

3.3.4 Dynamic auditing as an armed bandit problem

Optimal stopping problems are also intimately related to Bandit problems. Bandit problems are settings where a slot machine, referred to as a “bandit”, has $k + 1$ arms that yield different payoffs. All arms but one have payoffs governed by an unknown random variable, the last arm representing an outside option with known deterministic payoff that may be normalized to 0. A gambler must choose a limited number of arms to play at each period, and can update his beliefs about the payoffs after observing the outcome of his play. He thus faces an exploration/exploitation dilemma between choosing once and for all the best arm at the initial period, or explore by updating his choice given the accumulated experience. Hereafter, we place ourselves in an infinite horizon setting, but results are valid in the finite horizon problem.

3.3.4.1 One-armed bandit

We start with the most simple setting for a Bandit problem: the One-Armed Bandit. In this setting the player only faces one risky arm with returns $(X_1, X_2, \dots, X_n, \dots)$, and one safe arm with deterministic return λ at each period. Arms are characterized by their Gittins Index.

Definition 3.1 (Gittins Index for an arm). *For an arm characterized by the returns $(X_1, X_2, \dots, X_n, \dots)$, and N a stopping time, its Gittins Index is given by*

$$\Lambda(\delta) = \sup_{N \geq 1} \left\{ \frac{\mathbb{E}(\sum_1^N \delta^{j-1} X_j)}{\mathbb{E}(\sum_1^N \delta^{j-1})} \right\}.$$

Thus, the safe arm has a Gittins Index equal to $\lambda/(1 - \delta)$. The optimal policy is given by the Gittins Index Theorem for a One-Armed Bandit, whereby it is optimal to play the risky arm if $\Lambda(\delta) > \lambda/(1 - \delta)$. Otherwise, it is optimal to use the safe arm once and for all.

In our setting, the risky arm is auditing and the associated return are a function of the belief π . Its has a Gittins Index $\Lambda(\pi, \delta)$ that depends on π . The risky returns are given by $X_j = u(\tilde{\pi}_j, 1)$. The safe arm corresponds to not auditing, with $\lambda = 0$. When the Gittins Index of the risky arm is larger than 0, it is equivalent to the condition $\pi > \pi_\infty^*$.

Proposition 3.5.

$$\Lambda(\pi_1, \delta) > 0 \iff \pi_1 > \pi_\infty^*.$$

This is a direct consequence of the fact that $V_\infty(\pi_1) = \mathbb{E}[\sum_1^{\tau^*} \delta_{j-1} u(\tilde{\pi}_j, 1)]$ and $V_\infty(\pi_1) > 0$ for $\pi_1 > \pi_\infty^*$.

3.3.4.2 k-armed bandit

This characterization of our problem as a Bandit one allows us to answer the question of what happens with multiple auditees and capacity constrained auditing. Multi-Armed Bandits optimal strategies consist in Gittins Index Rules.

Definition 3.2 (Gittins Index Rule). *A decision rule that at each stage chooses an arm that has the highest Gittins index is called a **Gittins index rule**.*

The Gittins Index Theorem shows that such rules are optimal.

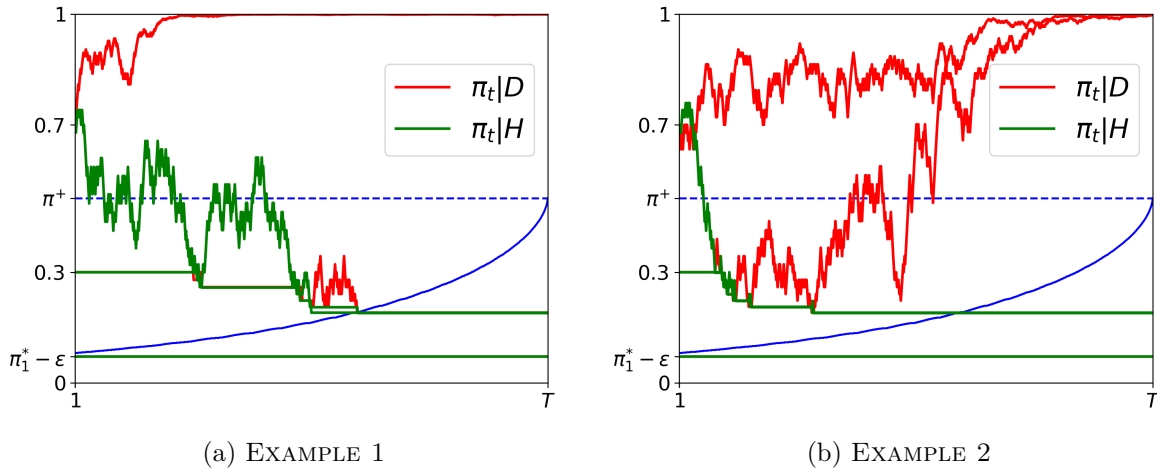
Theorem (The Gittins Index Theorem). *For a k -armed bandit problem with independent arms and geometric discounting, any Gittins index rule is optimal.*

We show in Proposition 3.6 that, in the context of our auditing problem, the Gittins Index Rule is equivalent to auditing arms by order of highest prior.

Proposition 3.6. *Given an initial prior π_1 , using a Gittins Index Rule is equivalent to auditing the highest priors in priority, conditional on the priors being larger than π_∞^* .*

$$\Lambda(\pi_1, \delta) > \Lambda(\pi'_1, \delta) \iff \pi_1 > \pi'_1 \quad (> \pi_\infty^*).$$

Figure 3.10 – 6-ARMED BANDIT WITH A CAPACITY CONSTRAINED AUDIT ($k = 2$)



Therefore, we can derive the optimal auditing strategy when there are multiple auditees and a capacity constrained audit. Let there be a population of auditees of size n . Let there also be a capacity constraint on auditing:

$$\sum_{i=0}^n x_{i,t} \leq k.$$

Then an optimal strategy is to audit by decreasing prior until the constraint is binding or the optimal threshold π_∞^* is reached. Figure 3.10 shows two examples of a 6-armed bandit when $k = 2$. The two auditees, one of each type, with the lowest initial prior at $\pi_1^* - \epsilon$ are never audited, because the outside option is more valuable since the beginning. In both examples, the type D starting at a prior of 0.7 is kept in the auditing pool. The other type D , starting at a lower prior of 0.3 is excluded from auditing in Figure 3.10a after a certain number of periods where it competes with the other two auditees of type H . In Figure 3.10b, this other type D , is kept in the auditing pool even if its prior decreases sharply at first. The two type H auditees are both excluded after a finite time.

3.4 The Extents of Learning

After showing how the possibility of learning through audits alters optimal auditing strategies, we now focus on different aspects of the learning process. First, in Section 3.4.1, we examine how much information the DM gets on average, i.e., how often he gets to know the truth, or at least how often he learns enough to take the optimal action. Second, in Section 3.4.2, we focus on the speed at which information is gathered, i.e., how fast we approach the truth.

3.4.1 How much can the DM learn?

We are hereafter interested in the question of whether the DM gets to learn the truth, i.e., how often the belief process converges to the true belief. In the Optimal Learning academic literature,¹⁰ this question is addressed through the concepts of complete and adequate learning (e.g., Easley and Kiefer (1988), Aghion et al. (1991)). Complete learning occurs when, with probability one, the agent acquires the true information about the state of the world. Adequate learning is a weaker version of complete learning, whereby it occurs when, “with probability one, the agent acquires enough information to allow him to obtain the true maximum payoff” (i.e., “learning everything worth knowing”).

Aghion et al. (1991) address the question of characterizing situations where adequate learning obtains or does not obtain. More specifically, they ask the question of whether adequate learning is the generic outcome. One situation where they find adequate learning to obtain is the undiscounted case, i.e., $\delta \rightarrow 1$.¹¹ This is exactly analogous to Statement iii in proposition 3.2: because of the absence of discounting, the optimal threshold tends to 0, which means systematic auditing, and therefore convergence of the belief process to its true value (Lemma 3.1). However, they argue that adequate learning is not the generic outcome, and define *partial learning* as occurring when the agent acquires adequate knowledge with probability strictly between 0 and 1. They suggest, through an informal discussion, that partial learning is the generic outcome.

Hereafter, we show that we can formally answer this question in our setting. With an infinite horizon, complete, adequate, and partial learning can be formulated through the distribution of the optimal stopping times. With a finite horizon, the comparison is not possible as we cannot talk about the convergence of beliefs.

¹⁰In this literature, the horizon is usually infinite.

¹¹Other exhibited situations where adequate learning occurs include smoothness and quasi-concavity, or analyticity of the payoff function

Infinite horizon In the infinite horizon setting, adequate learning occurs if, in the case of a type H , the belief converges to a limit for which there is no auditing, or if, in the case of a type D , the belief converges to a limit for which there is auditing. In other words, conditional on type H , adequate learning occurs when $\tau^* < \infty$, and the belief process drops to a value π^{lim} below π_∞^* . The belief therefore converges to π^{lim} since auditing stops. Conditional on type D , adequate learning occurs when $\tau^* = \infty$, auditing never stops, and the belief process converges to 1. Definition 3.3 formalizes this idea and Proposition 3.7 answers the question of genericity.

Definition 3.3 (Adequate learning in infinite horizon). *Let τ^* be the stopping time $\tau^* = \inf(t \in \mathbb{N}^* | \pi_t < \pi_\infty^*)$. For an initial prior π_1 , adequate learning occurs with probability*

$$\begin{aligned} \mathbb{P}(AL|\pi_1) &= \pi_1 \mathbb{P}(\tau^* = \infty | D, \pi_1) + (1 - \pi_1)(1 - \mathbb{P}(\tau^* = \infty | H, \pi_1)), \\ &= \pi_1 \mathbb{P}(\tau^* = \infty | D, \pi_1) + (1 - \pi_1). \end{aligned} \quad (3.4)$$

Proposition 3.7. *In this setting, partial learning is the generic outcome for all initial priors π_1*

$$0 < \mathbb{P}(AL|\pi_1) < 1, \quad \forall \pi_1 \in (0, 1). \quad (3.5)$$

In particular

$$\mathbb{P}(AL|\pi_1 < \pi_\infty^*) = (1 - \pi_1). \quad (3.6)$$

Adequate learning is never obtained because it is too strong a condition. Adequacy is always obtained conditional on type H , but never for type D . When the initial prior is below the optimal threshold, adequate learning occurs if the auditee is of type H , i.e., with probability $(1 - \pi_1)$, but not for type D , as he is never audited. This means $\mathbb{P}(\tau^* = \infty | D, \pi_1) = 0$ and explains equation (3.6). When the initial prior is larger than π_∞^* , there is a non null probability, conditional on type D , for the belief process to hit the optimal threshold, and $\mathbb{P}(\tau^* = \infty | D, \pi_1) < 1$. Example 3.2 below provides a special case where $\mathbb{P}(\tau^* = \infty | D, \pi_1)$ is computable.

Example 3.2 (Infinite horizon Gambler's Ruin). *When $p_D + p_H = 1$, the updating functions are symmetrical ($\varphi_{AI} = \varphi_{AV}^{-1}$). Given an infinite horizon threshold π_∞^* and an initial prior π_1 , the updating functions define a partition of $(0, 1) = \bigcup_{k \in \mathbb{Z}} \Phi_k = \bigcup_{k \in \mathbb{Z}} (\varphi_{AI}^k(\pi_1), \varphi_{AI}^{k+1}(\pi_1)]$, and there exists a k^* such that $\pi_\infty^* \in \Phi_{k^*}$. Then the problem is analogous to a Gambler's Ruin problem with*

a biased toss, where the player wins with probability $p_D > \frac{1}{2}$ and has initial fortune k^* . Thus

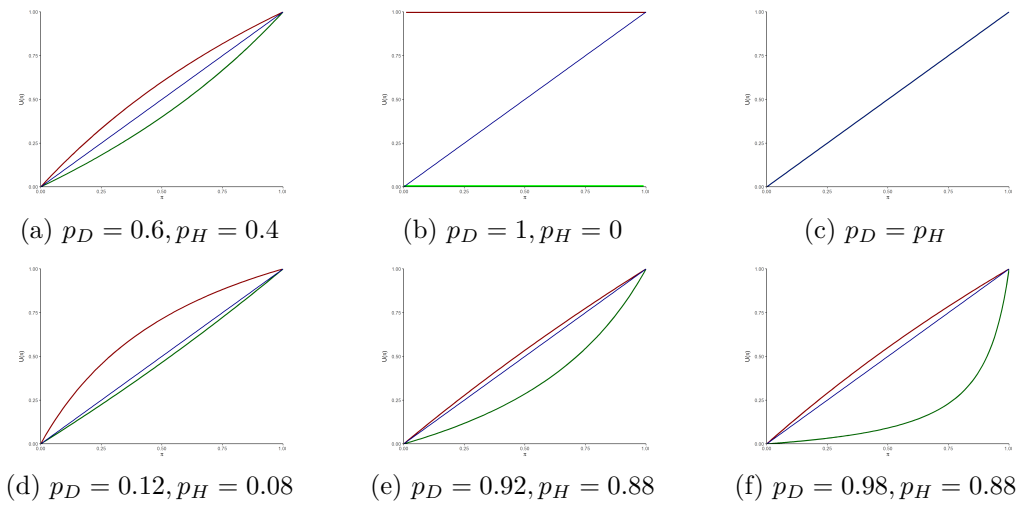
$$\mathbb{P}(\tau^* = \infty | D, \pi_1) = \begin{cases} 0 & \text{if } k^* \geq 0, \\ 1 - \left(\frac{p_D}{1-p_D}\right)^{k^*} & \text{if } k^* < 0, \end{cases}$$

and, when $\pi_1 > \pi_\infty^*$, partial learning occurs, with adequate learning happening with probability

$$\mathbb{P}(AL | \pi_1) = \pi_1 \left(1 - \left(\frac{p_D}{1-p_D}\right)^{k^*}\right) + (1 - \pi_1).$$

3.4.2 How fast can the DM learn?

Figure 3.11 – UPDATING FUNCTIONS FOR DIFFERENT PARAMETERS



Lemma 3.1 shows that the belief process is a martingale and converges to the true belief. Now, we are interested in how fast information is accumulated, and how this speed of information acquisition is related to the parameters of the model.

Figure 4.3b shows the updating functions $\{\varphi_{AI}, \varphi_{AV}\}$ for different values of the parameters (p_D, p_H) . The first intuition is that the more polarized the behaviors of types H and D , i.e., the larger $p_D - p_H$, the more informative the audit. The extreme case $p_D - p_H = 1$ is illustrated in Figure 3.11b. In this case, auditing is perfectly informative, as observing an invalid claim implies that the auditee is type D , while observing a valid claim implies he is of type H . On the contrary, when $p_D = p_H$ (Figure 3.11c), types behave the same way and auditing reveals nothing about the type. The second intuition is related to the informativeness of extreme behaviors. i.e., when for a fixed difference $p_D - p_H$, either p_D is close to 1 or p_H is close to 0. If $p_D \approx 1$, observing a valid claim is a strong sign that the auditee is of type H , as shown in Figure 3.11f. If $p_H \approx 0$, observing an invalid claim is a strong sign that the auditee is of type D .

3.4.2.1 Doob-Meyer decomposition

In proposition 3.8, we use the Doob-Meyer decomposition for submartingales and supermartingales to decompose the belief processes under both states of the world H and D . This decomposition allows us to identify the magnitude of information acquisition.

Proposition 3.8. *[Information acquisition] Under states of the world H and D , we can decompose the stochastic process of updated beliefs thanks to the Doob-Meyer decomposition.*

i Under \mathbb{P}_H , $\tilde{\pi}_t = M_t^H + P_t^H$ where M_t^H is a martingale and P_t^H is a decreasing and predictable process.

$$ii \ P_t^H = -(p_D - p_H) \sum_{i=1}^{t-1} \tilde{\pi}_i \left[\varphi_{AI}(\tilde{\pi}_i) - \varphi_{AV}(\tilde{\pi}_i) \right] < 0$$

iii Under \mathbb{P}_D , $\tilde{\pi}_t = M_t^D + P_t^D$ where M_t^D is a martingale and P_t^D is an increasing and predictable process.

$$iv \ P_t^D = (p_D - p_H) \sum_{i=1}^{t-1} (1 - \tilde{\pi}_i) \left[\varphi_{AI}(\tilde{\pi}_i) - \varphi_{AV}(\tilde{\pi}_i) \right] > 0$$

The predictable process P_t^i represents the average gain in information under the true state $i \in \{H, D\}$

$$\mathbb{E}_i[\tilde{\pi}_t | \pi_1] = \pi_1 + \mathbb{E}_i[P_t^i | \pi_1].$$

3.4.2.2 Informativeness and exploration

The decomposition in Proposition 3.8 allows us to quantify the amount of information obtained at each stage and to understand how it affects the optimal level of exploration. From the expression of P_t , we can see that the difference $\beta = p_D - p_H$ plays an important role. This is rather intuitive in that a larger difference means more polarized behavior between the Honest and Dishonest behaviors. Proposition 3.9 formalizes the idea that the more polarized the behavior, the larger $|P_t^H|$ and $|P_t^D|$, and the more information is gathered on average.

Proposition 3.9. *At any time t , the predictable process P_t^H is increasing in p_H and decreasing in p_D*

$$\frac{\partial P_t^H}{\partial p_H} > 0 \quad \text{and} \quad \frac{\partial P_t^H}{\partial p_D} < 0,$$

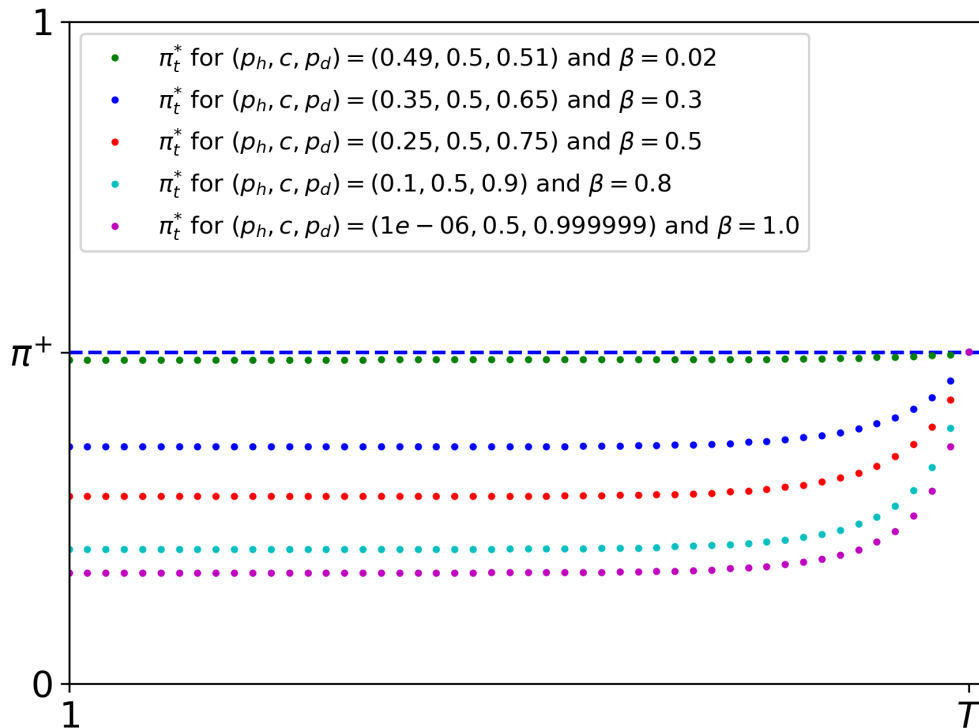
and the predictable process P_t^D is increasing in p_D and decreasing in p_H

$$\frac{\partial P_t^D}{\partial p_D} > 0 \quad \text{and} \quad \frac{\partial P_t^D}{\partial p_H} < 0.$$

Behaviors are more polarized when p_D increases and/or p_H decreases. When it is the case, P_t^H is smaller, and since $P_t^H < 0$, the average decrease in the belief $|P_t^H|$ is larger, and beliefs converge towards the true belief 0 faster. The reasoning is analogous for type D as, when behaviors are more polarized, P_t^D is larger. As $P_t^D > 0$, the average increase in the belief P_t^D is larger and beliefs converge towards the true belief 1 faster.

The impact of a more efficient learning, in the sense that belief processes converge faster, is such that optimal thresholds are lower, and more exploration takes place under optimal auditing at each period. Figure 3.13 illustrates different optimal threshold sequences for different levels of polarization of behaviors $\beta = p_D - p_H$.

Figure 3.13 – POLARIZED AND EXTREME FRAUD BEHAVIORS INDUCE MORE EXPLORATION



3.5 Conclusion

This article investigates the role of learning in the context of fraud inspection when inspector/inspectee interactions are repeated. Auditing acts as a learning mechanism that isolates the

dishonest (worth auditing) types from the honest (unprofitable to audit) types.

Our first set of results shows how future auditing proceeds are always increasing in the current auditing efforts, but current proceeds may be negative. Optimal auditing relies then on an exploration/exploitation trade-off consisting in balancing auditing-diminished present proceeds and auditing-enhanced discounted future proceeds. Because learning-generated advantages reverberate through all future periods, the number of remaining periods positively impact auditing efforts. As a weighting of future proceeds, low discounting encourages learning through auditing and complete exploration occurs in an infinite horizon when the auditor is indifferent between present and future proceeds.

Our second set of results examines the particularities of our learning process: we first show that, in our setting, partial learning is the generic outcome. Second, we identify the average change in the belief process as the predictable component of a Doob-Meyer decomposition of the belief process in each state of the world. This average change is larger when dishonest and honest types behave very differently, as the separating power of auditing is then stronger.

3.6 Appendix

3.6.1 Proofs

3.6.1.1 Proof of Lemma 3.1

Proof. At period $t + 1$, with prior π_t and auditing decision x_t

$$\begin{aligned} \mathbb{E}_{\mathbb{P}}[\tilde{\pi}_{t+1} | \mathcal{F}_t, x_t] &= x_t[\bar{p}(\pi)\varphi_{AI}(\pi_t) + (1 - \bar{p}(\pi))\varphi_{AV}(\pi_t)] + (1 - x_t)[\varphi_{NA}(\pi_t)] \\ &= x_t[p_D\pi_t + (1 - p_D)\pi_t] + (1 - x_t)\pi_t \\ &= \pi_t \end{aligned}$$

$$\begin{aligned} \mathbb{E}_D[\tilde{\pi}_{t+1} | \mathcal{F}_t, x_t] &= x_t[p_D\varphi_{AI}(\pi_t) + (1 - p_D)\varphi_{AV}(\pi_t)] + (1 - x_t)[\varphi_{NA}(\pi_t)] \\ &= x_t[\bar{p}(\pi)\varphi_{AI}(\pi_t) + (1 - \bar{p}(\pi))\varphi_{AV}(\pi_t) + (p_D - \bar{p}(\pi))(\varphi_{AI}(\pi_t) - \varphi_{AV}(\pi_t))] + (1 - x_t)\pi_t \\ &= \pi_t + x_t \underbrace{[(p_D - \bar{p}(\pi))]}_{\geq 0} \underbrace{(\varphi_{AI}(\pi_t) - \varphi_{AV}(\pi_t))}_{\geq 0} \\ &\geq \pi_t \end{aligned}$$

$$\begin{aligned} \mathbb{E}_H[\tilde{\pi}_{t+1} | \mathcal{F}_t, x_t] &= x_t[p_H\varphi_{AI}(\pi_t) + (1 - p_H)\varphi_{AV}(\pi_t)] + (1 - x_t)[\varphi_{NA}(\pi_t)] \\ &= x_t[\bar{p}(\pi)\varphi_{AI}(\pi_t) + (1 - \bar{p}(\pi))\varphi_{AV}(\pi_t) + (p_H - \bar{p}(\pi))(\varphi_{AI}(\pi_t) - \varphi_{AV}(\pi_t))] + (1 - x_t)\pi_t \\ &= \pi_t + x_t \underbrace{[(p_H - \bar{p}(\pi))]}_{\leq 0} \underbrace{(\varphi_{AI}(\pi_t) - \varphi_{AV}(\pi_t))}_{\geq 0} \\ &\leq \pi_t \end{aligned}$$

Finally, the convergence of the belief process to the true belief is a direct consequence of the Martingale Convergence Theorem for bounded martingales. \square

3.6.1.2 Proof of Corollary 3.1

Proof. This is a direct consequence of the martingale property. \square

3.6.1.3 Proof of Lemma 3.2

This is straightforward as

- $\bar{p}(\pi) - c$ is strictly increasing in π , strictly negative for $\pi = 0$ and strictly positive for $\pi = 1$,

- π^+ is then uniquely defined,
- $\text{sign}\left(\frac{\partial\Omega_T}{\partial x}\right) = \text{sign}(\pi - \pi^+)$,
- then $x_T^* = \mathbb{1}\{\pi \geq \pi^+\}$.

3.6.1.4 Proof of Lemma 3.3

Proof. We proceed by backward induction for $t \in \{1, \dots, T\}$.

At the final stage $t = T$: The objective function

$$\Omega_{T,T}(\pi, x) = u(\pi, x) = (\bar{p}(\pi) - c)x,$$

is trivially (piece-wise) linear in π , linear in x . Its derivative with regards to x

$$\frac{\partial\Omega_{T,T}}{\partial x}\Big|_{\pi} = \bar{p}(\pi) - c = (p_D - p_H)\pi + p_H - c,$$

is increasing in π , is equal to $p_D - c > 0$ for $\pi = 1$ and to $p_H - c < 0$ for $\pi = 0$. It implies the existence of a threshold $\pi_{T,T}^* \in (0, 1)$ such that

$$\frac{\partial\Omega_{T,T}}{\partial x}\Big|_{\pi_{T,T}^*} = 0,$$

and

$$x_{T,T}^*(\pi) \begin{cases} = 0 & \text{if } \pi < \pi_{T,T}^*, \\ \in (0, 1) & \text{if } \pi = \pi_{T,T}^*, \\ = 1 & \text{if } \pi > \pi_{T,T}^*. \end{cases}$$

Then, the value function

$$V_{T,T}(\pi) = \max(0, \bar{p}(\pi) - c),$$

is (piece-wise) linear in π and convex as the maximum of the finite family of linear functions indexed by \mathcal{I}_T

$$\mathcal{I}_T = \{0, \bar{p}(\pi) - c\}.$$

At any stage $t \in \{1, \dots, T - 1\}$: The objective function

$$\begin{aligned}\Omega_{t,T}(\pi, x) &= u(\pi, x) + \delta \mathbb{E}_t [V_{t+1,T}(\tilde{\pi}_{t+1})], \\ &= (\bar{p}(\pi) - c)x + \delta \left[x \left(\bar{p}(\pi) V_{t+1,T}(\varphi_{AI}(\pi)) + (1 - \bar{p}(\pi)) V_{t+1,T}(\varphi_{AV}(\pi)) \right) + (1 - x) V_{t+1,T}(\pi) \right],\end{aligned}$$

is trivially linear in x as x appears only in the transition probabilities, while $V_{t+1,T}$ is independent of x . In addition, by induction, $V_{t+1,T}$ is piece-wise linear in π . It implies, for every $\pi \in (0, 1)$, the existence of a couple $(\alpha, \beta) \in \mathbb{R}^2$ such that

$$V_{t+1,T}(\pi) = \alpha\pi + \beta.$$

Therefore,

$$\begin{aligned}\bar{p}(\pi) V_{t+1,T}(\varphi_{AI}(\pi)) &= \bar{p}(\pi) (\alpha \varphi_{AI}(\pi) + \beta) = \alpha p_D \pi + \bar{p}(\pi) \beta, \\ \bar{p}(\pi) V_{t+1,T}(\varphi_{AV}(\pi)) &= (1 - \bar{p}(\pi)) (\alpha' \varphi_{AV}(\pi) + \beta') = \alpha' (1 - p_D) \pi + (1 - \bar{p}(\pi)) \beta',\end{aligned}$$

and $\Omega_{t,T}$ is piece-wise linear in π . It is also continuous in π as the sum of functions continuous in π . Its derivative with regards to x verifies

$$\begin{aligned}\left. \frac{\partial \Omega_{t,T}}{\partial x} \right|_{\pi} &= \bar{p}(\pi) - c + \delta \left[\bar{p}(\pi) V_{t+1,T}(\varphi_{AI}(\pi)) + (1 - \bar{p}(\pi)) V_{t+1,T}(\varphi_{AV}(\pi)) - V_{t+1,T}(\pi) \right], \\ \left. \frac{\partial \Omega_{t,T}}{\partial x} \right|_{\pi=0} &= p_H - c < 0, \\ \left. \frac{\partial \Omega_{t,T}}{\partial x} \right|_{\pi=1} &= p_D - c > 0.\end{aligned}$$

In addition, a second derivative with regards to π

$$\begin{aligned}\left. \frac{\partial \Omega_{t,T}}{\partial x \partial \pi} \right|_{\pi} &= p_D - p_H + \delta \left[\alpha (\varphi_{AI}(\pi)) p_D + \alpha (\varphi_{AV}(\pi)) (1 - p_D) - \alpha(\pi) \right], \\ &> 0,\end{aligned}$$

is strictly positive,¹² thus defining the threshold $\pi_{t,T}^*$. Since the optimal action is either to audit or not to audit, the value function can be expressed as

$$\begin{aligned} V_{t,T}(\pi) &= \max(\Omega_{t,T}(\pi, 0), \Omega_{t,T}(\pi, 1)) \\ &= \max(0, \bar{p}(\pi) - c + \delta V_{t+1}(\tilde{\pi}_{t+1})) \\ &= \max(0, \bar{p}(\pi) - c + \delta[\bar{p}(\pi)V_{t+1}(\varphi_{AI}(\pi)) + (1 - \bar{p}(\pi))V_{t+1}(\varphi_{AV}(\pi))]). \end{aligned}$$

Then, for any element $fl(\pi)$ of the linear functions family \mathcal{I}_{t+1} such that $V_{t+1}(\pi) = \max_{fl \in \mathcal{I}_{t+1}} fl(\pi)$, reasoning as previously, $\bar{p}(\pi)fl_i(\varphi_{AI}(\pi))$ and $(1 - \bar{p}(\pi))fl_i(\varphi_{AV}(\pi))$ are linear in π , therefore $\bar{p}(\pi)V_{t+1}(\varphi_{AI}(\pi))$ and $(1 - \bar{p}(\pi))V_{t+1}(\varphi_{AV}(\pi))$ are maximums of families of linear functions. Finally, using usual properties of sums or maximums, scalar multiplication of maximums and maximums of maximums, there exists a family \mathcal{I}_t of linear functions such that

$$V_t(\pi) = \max_{fl \in \mathcal{I}_t} fl(\pi).$$

As the maximum of a finite family of linear functions, $V_t(\pi)$ is continuous, piece-wise linear and convex in π . □

3.6.1.5 Proof of Lemma 3.4

Proof. This is a direct consequence of the martingale property, the convexity of V_{t+1} in π and Jensen's inequality:

$$\begin{aligned} \frac{\partial}{\partial x_t} \mathbb{E}_t[V_{t+1}(\tilde{\pi}_{t+1})|x_t] &= \bar{p}(\pi_t)V_{t+1}(\varphi_{AI}(\pi_t)) + (1 - \bar{p}(\pi_t))V_{t+1}(\varphi_{AV}(\pi_t)) - V_{t+1}(\pi_t) \\ &> V_{t+1}(\bar{p}(\pi_t)\varphi_{AI}(\pi_t) + (1 - \bar{p}(\pi_t))\varphi_{AV}(\pi_t) - \pi_t) \\ &\geq 0. \end{aligned}$$

The positivity is straightforward as

$$V_t(\pi) \geq \Omega_t(\pi, x_t = 0) = \delta V_{t+1}(\pi) \geq \dots \geq \delta^{T-t} V_T(\pi) \geq 0.$$

□

¹²Because $p_D - p_H = \max \alpha$

3.6.1.6 Proof of Proposition 3.1

Proof. We proceed by backward induction.

Initialization Stage $t = T - 1$: The objective function verifies

$$\left. \frac{\partial \Omega_{T-1,T}}{\partial x} \right|_{\pi} = \bar{p}(\pi) - c + \delta \left[\bar{p}(\pi) V_{T,T}(\varphi_{AI}(\pi)) + (1 - \bar{p}(\pi)) V_{T,T}(\varphi_{AV}(\pi)) - V_{T,T}(\pi) \right],$$

and, since $V_{T,T}(\pi) = 0$ below $\pi_{T,T}^*$ and $\varphi_{AV}(\pi) < \pi < \varphi_{AI}(\pi)$,

$$\left. \frac{\partial \Omega_{T-1,T}}{\partial x} \right|_{\pi_{T,T}^*} = \delta \bar{p}(\pi_{T,T}^*) V_{T,T}(\varphi_{AI}(\pi_{T,T}^*)) > 0.$$

As $\pi_{T-1,T}^*$ is defined by $\left. \frac{\partial \Omega_{T-1,T}}{\partial x} \right|_{\pi_{T-1,T}^*} = 0$ and $\frac{\partial \Omega_{T-1,T}}{\partial x}$ is strictly increasing in π , it implies

$$\pi_{T-1,T}^* < \pi_{T,T}^*.$$

At any stage $t \in \{1, \dots, T - 2\}$: The objective function verifies

$$\left. \frac{\partial \Omega_{t,T}}{\partial x} \right|_{\pi} = \bar{p}(\pi) - c + \delta \left[\bar{p}(\pi) V_{t+1,T}(\varphi_{AI}(\pi)) + (1 - \bar{p}(\pi)) V_{t+1,T}(\varphi_{AV}(\pi)) - V_{t+1,T}(\pi) \right],$$

and, since $V_{t+1,T}(\pi) = 0$ below $\pi_{t+1,T}^*$ and $\varphi_{AV}(\pi) < \pi < \varphi_{AI}(\pi)$,

$$\left. \frac{\partial \Omega_{t+1,T}}{\partial x} \right|_{\pi_{t+1,T}^*} = 0 = \bar{p}(\pi_{t+1,T}^*) - c + \delta \left[\bar{p}(\pi_{t+1,T}^*) V_{t+1,T}(\varphi_{AI}(\pi_{t+1,T}^*)) \right],$$

thus, by subtraction,

$$\left. \frac{\partial \Omega_{t,T}}{\partial x} \right|_{\pi_{t+1,T}^*} = \delta \left[\bar{p}(\pi_{t+1,T}^*) (V_{t+1,T} - V_{t+2,T})(\varphi_{AI}(\pi_{t+1,T}^*)) \right],$$

Since at prior π and time $t + 1$, playing the sequence $(x_{t+1,T}^*, x_{t+2,T}^*, \dots, x_{T,T}^*)$ is optimal and yields $V_{t+1,T}(\pi)$, it dominates any other sequence, in particular $(x_{t+2,T}^*, x_{t+3,T}^*, \dots, x_{T,T}^*, 0)$, which yields $V_{t+2,T}(\pi)$. Hence

$$V_{t+2,T}(\varphi_{AI}(\pi_{t+1,T}^*)) \leq V_{t+1,T}(\varphi_{AI}(\pi_{t+1,T}^*)).$$

In addition, since $\varphi_{AI}(\pi_{t+1,T}^*) > \pi_{t+2,T}^*$, then $A^k(\pi_{t+1,T}^*) > \pi_{t+k+1,T}^*$ for $k \in \{1, T - t - 1\}$, and the corresponding belief transitions happen with strictly positive probability $\in (p_H^{T-t}, p_D^{T-t})$. It im-

plies that the sequence $(x_{t+2,T}^*, x_{t+3,T}^*, \dots, x_{T,T}^*, x_{T,T}^*)$ strictly dominates $(x_{t+2,T}^*, x_{t+3,T}^*, \dots, x_{T,T}^*, 0)$.

Therefore

$$V_{t+2,T}(\varphi_{AI}(\pi_{t+1,T}^*)) < V_{t+1,T}(\varphi_{AI}(\pi_{t+1,T}^*)).$$

In the end

$$\left. \frac{\partial \Omega_{t,T}}{\partial x} \right|_{\pi_{t+1,T}^*} > 0 = \left. \frac{\partial \Omega_{t,T}}{\partial x} \right|_{\pi_{t,T}^*}.$$

Since $\frac{\partial \Omega_{t,T}}{\partial x}$ is a strictly increasing function of π ,

$$\pi_{t,T}^* < \pi_{t+1,T}^*.$$

This proves statement [i](#). Statement [ii](#) is a direct consequence by backward induction. Statement [iii](#) is a result of both these statements as

$$\pi_{1,T}^* = \pi_{2,T+1}^* > \pi_{1,T+1}^*.$$

As for the impact of δ , when $\delta \rightarrow 0$, the result is straightforward as the problem comes down to the myopic one. A simple backward induction from stage $t = T - 1$ shows that

$$\frac{\partial}{\partial \delta} V_{t,T}(\pi, \delta) > 0, \quad \frac{\partial}{\partial x \partial \delta} \Omega_{t,T} > 0,$$

implying that, since $\pi_{t,T}^*(\delta)$ solves for $\left. \frac{\partial}{\partial x} \Omega_{t,T} \right|_{\pi, \delta} = 0$ and $\frac{\partial \Omega}{\partial x}$ is decreasing in π ,

$$\frac{\partial \pi_{t,T}^*}{\partial \delta} < 0.$$

□

3.6.1.7 Proof of Proposition [3.2](#)

Proof. Statement [i](#) is a direct consequence of $(\pi_{t,T}^*)_{T/in\mathbb{N}^*}$ being decreasing and bounded below by 0. Statement [ii](#) comes from the fact that at each stage the value function is the maximum of a finite family of linear functions.

As for Statement [iii](#), assume that it is not true, i.e., $\lim_{\delta \rightarrow 1} \pi_{\infty}^*(\delta) = \underline{\pi} > 0$. Then, for $\pi < \underline{\pi}$, $V_{\infty}(\pi) = 0$. Consider the profits of the auditor under strategy $(x_t = 1)_{1 \leq t \leq \bar{t}}$. Because of the

linearity of u and the martingale property of beliefs

$$\mathbb{E}_1 \left[\sum_{i=1}^{\bar{t}} \delta^{i-1} u(\tilde{\pi}_i, 1) \right] = \frac{1 - \delta^{\bar{t}}}{1 - \delta} u(\pi, 1).$$

With probability π , the auditee is initially believed to be of type D , and because of the convergence of beliefs to the true state, there exists $r \in (0, 1)$ and a time t_r such that

$$\mathbb{P}(\tilde{\pi}_{t_r} > \pi^+ + \varepsilon | D, (x_t)_{1 \leq t \leq t_r} = 1) > r.$$

Then, consider the alternative strategy $(x_t^{al})_t$ defined as

$$x_t^{al} = \mathbb{1}_{\{t \leq t_r\}} + \mathbb{1}_{\{t > t_r\}} \mathbb{1}_{\{\pi_{t_r} > \pi^+ + \varepsilon\}},$$

with corresponding expected profits higher than

$$\frac{1 - \delta^{t_r}}{1 - \delta} \underbrace{u(\pi, 1)}_{<0} + \pi r \frac{\delta^{t_r}}{1 - \delta} \underbrace{u(\pi^+ + \varepsilon, 1)}_{>0}.$$

When δ tends to one, the left hand side term is finite, while the right hand side term goes to infinity. This last result contradicts the fact that $V_\infty(\pi) = 0$. □

3.6.1.8 Proof of Proposition 3.3

Proof. First, τ^* is indeed a stopping time as it is defined by the process $\langle \pi_t \rangle$ that is by definition adapted to the filtration $(\mathcal{F}_t)_{t \in \{1, \dots, T\}}$. In addition,

$$\forall t < \tau^*, \quad \tilde{\pi}_t \geq \pi_t^*, \tag{3.7}$$

$$\forall t \geq \tau^*, \quad \tilde{\pi}_t = \tilde{\pi}_{\tau^*} < \pi_{\tau^*}^* < \pi_t^*. \tag{3.8}$$

Equation (3.7) is true because otherwise, by definition of τ^* as a minimum of times at which the previous inequality is verified, $\tau^* \leq t$. In equation (3.8), the first equality relies on the fact that once audit stops, i.e., at τ^* , priors remain constant, the following inequality stems from the definition of τ^* , and finally the last inequality stems from the fact that the sequence $(\pi_t^*)_t$ is increasing, from Proposition 3.1.

Then rewriting the value function yields

$$\begin{aligned}
 V_{1,T}(\pi_1, \delta) &= \mathbb{E} \left[\sum_{j=1}^T \delta^{j-1} u(\tilde{\pi}_j, \mathbb{1}\{\tilde{\pi}_j \geq \pi_j^*\}) \right] \\
 &= \mathbb{E} \left[\sum_{j=1}^T \delta^{j-1} u(\tilde{\pi}_j, 1) \mathbb{1}\{\tilde{\pi}_j \geq \pi_j^*\} \right] \\
 &= \mathbb{E} \left[\sum_{j=1}^{\tau^*} \delta^{j-1} u(\tilde{\pi}_j, 1) \right], \quad \text{from Eq. (3.7) and (3.8)} \\
 &= W_{1,\tau^*}(\pi_1, \delta).
 \end{aligned}$$

Finally, τ^* maximizes W because otherwise, there would be another stopping time based auditing strategy that would maximize W and yield auditing proceeds strictly higher than $V_{1,T}(\pi_1, \delta)$, which would be contradictory. \square

3.6.1.9 Proof of Proposition 3.4

Proof. For a given π_1 and the corresponding optimal stopping rule τ^*

$$\begin{aligned}
 W_{1,\tau^*}(\pi_1, \delta) &= \mathbb{E} \left[\sum_{j=1}^{\tau^*-1} \delta^{j-1} u(\tilde{\pi}_j, 1) \right] \\
 &= \sum_{\omega} \mathbb{P}(\omega) \sum_{j=1}^{\tau^*(\omega)-1} \delta^{j-1} u(\pi_j(\omega), 1).
 \end{aligned}$$

Using the same reasoning as in the proof of Lemma 3.3, $u(\pi_j(\omega), 1)$ is a linear function of the initial belief π_1 , $u(\pi_j(\omega), 1) = \alpha_j(\omega)\pi_1 + \beta_j(\omega)$. Thus

$$\frac{\partial W_{1,\tau^*}}{\partial \pi_1} = \sum_{\omega} \mathbb{P}(\omega) \sum_{j=1}^{\tau^*(\omega)-1} \delta^{j-1} \alpha_j(\omega).$$

This is the slope of the value function at π_1 . In addition, notice that for $\pi'_1 < \pi_1$ such that they are not in the same Ψ_m

$$\tau^*(\pi'_1, \omega) \leq \tau^*(\pi_1, \omega)$$

and there exists at least one trajectory ω' such that the inequality is strict. Otherwise, by definition of the intervals Ψ , π'_1 and π_1 would be in the same Ψ_m . Then

$$\left. \frac{\partial W_{1,\tau^*}}{\partial \pi_1} \right|_{\pi'_1} < \left. \frac{\partial W_{1,\tau^*}}{\partial \pi_1} \right|_{\pi_1}.$$

This last inequality shows that the linear functions slopes are increasing in π , as shown in Figure 3.5e. □

3.6.1.10 Proof of Proposition 3.6

Proof. We first show the direct implication of the equivalence. Denoting $\tau^*(\pi)$ and $\tau^*(\pi')$ the optimal stopping times for each prior

$$\begin{aligned}
 \Lambda(\pi, \delta) > \Lambda(\pi', \delta) &\Rightarrow \frac{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi)} \delta^{j-1} u(\tilde{\pi}_j, 1)]}{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi)} \delta^{j-1}]} > \frac{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi')} \delta^{j-1} u(\tilde{\pi}'_j, 1)]}{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi')} \delta^{j-1}]} \\
 &\Rightarrow \frac{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi)} \delta^{j-1} u(\tilde{\pi}_j, 1)]}{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi)} \delta^{j-1}]} > \frac{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi)} \delta^{j-1} u(\tilde{\pi}'_j, 1)]}{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi)} \delta^{j-1}]} \\
 &\Rightarrow \mathbb{E}\left[\sum_{j=1}^{\tau^*(\pi)} \delta^{j-1} u(\tilde{\pi}_j, 1)\right] > \mathbb{E}\left[\sum_{j=1}^{\tau^*(\pi)} \delta^{j-1} u(\tilde{\pi}'_j, 1)\right] \\
 &\Rightarrow \pi > \pi'.
 \end{aligned}$$

Then, for the indirect implication

$$\begin{aligned}
 \pi > \pi' &\Rightarrow \forall \omega, u(\pi_j(\omega), 1) > u(\pi'_j(\omega), 1) \\
 &\Rightarrow \frac{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi')} \delta^{j-1} u(\tilde{\pi}_j, 1)]}{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi')} \delta^{j-1}]} > \frac{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi')} \delta^{j-1} u(\tilde{\pi}'_j, 1)]}{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi')} \delta^{j-1}]} = \Lambda(\pi', \delta) \\
 &\Rightarrow \Lambda(\pi, \delta) \geq \frac{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi')} \delta^{j-1} u(\tilde{\pi}_j, 1)]}{\mathbb{E}[\sum_{j=1}^{\tau^*(\pi')} \delta^{j-1}]} > \Lambda(\pi', \delta).
 \end{aligned}$$

□

3.6.1.11 Proof of Proposition 3.7

Proof. Because of the $(1 - \pi_1)$ component, it is sufficient to show that $\mathbb{P}(\tau^* = \infty | D, \pi_1) < 1$. This is straightforward as there exists some $l \in \mathbb{N}$ such that $\varphi_{AV}^l(\pi_1) < \pi_\infty^* \leq \varphi_{AV}^{l-1}(\pi_1)$.¹³ This is the prior obtained from auditing as long as $\pi_1 \geq \pi_\infty^*$ and observing valid claims every time, which happens with probability $(1 - p_D)^l$. Therefore, $\mathbb{P}(\tau^* = \infty | D, \pi_1) \leq 1 - (1 - p_D)^l$. □

3.6.1.12 Proof of Proposition 3.8

Proof. Let $\langle I_t \rangle$ be the Bernoulli process representing the fraud outcome at period. Then $\mathbb{P}_H(I_t = 1) = p_H$ and $\mathbb{P}_D(I_t = 1) = p_D$. The updated prior under a given state of the world H

¹³If $\pi_1 < \pi_\infty^*$, then $l = 0$.

or D can be written as

$$\begin{aligned}\tilde{\pi}_{t+1} &= I_{t+1}\varphi_{AI}(\tilde{\pi}_t) + (1 - I_{t+1})\varphi_{AV}(\tilde{\pi}_t) \\ &= \tilde{\pi}_t + \left(I_{t+1} - \bar{p}(\tilde{\pi}_t)\right) \left[\varphi_{AI}(\tilde{\pi}_t) - \varphi_{AV}(\tilde{\pi}_t)\right]\end{aligned}$$

In particular

$$\begin{aligned}\mathbb{E}_D[\tilde{\pi}_{t+1}|\mathcal{F}_t] &= \tilde{\pi}_t + \left(p_D - \bar{p}(\tilde{\pi}_t)\right) \left[\varphi_{AI}(\tilde{\pi}_t) - \varphi_{AV}(\tilde{\pi}_t)\right] \\ &= \tilde{\pi}_t + \left(1 - \tilde{\pi}_t\right) \left(p_D - p_H\right) \left[\varphi_{AI}(\tilde{\pi}_t) - \varphi_{AV}(\tilde{\pi}_t)\right]\end{aligned}$$

and

$$\begin{aligned}\mathbb{E}_H[\tilde{\pi}_{t+1}|\mathcal{F}_t] &= \tilde{\pi}_t + \left(p_H - \bar{p}(\tilde{\pi}_t)\right) \left[\varphi_{AI}(\tilde{\pi}_t) - \varphi_{AV}(\tilde{\pi}_t)\right] \\ &= \tilde{\pi}_t - \tilde{\pi}_t \left(p_D - p_H\right) \left[\varphi_{AI}(\tilde{\pi}_t) - \varphi_{AV}(\tilde{\pi}_t)\right]\end{aligned}$$

One decomposition into a martingale and a predictable process is given, under state of the world D , by

$$\tilde{\pi}_t = M_t^D + P_t^D$$

M_t^D is a martingale and P_t^D is predictable and increasing

$$P_t^D = (p_D - p_H) \sum_{i=1}^{t-1} \left(1 - \tilde{\pi}_i\right) \left[\varphi_{AI}(\tilde{\pi}_i) - \varphi_{AV}(\tilde{\pi}_i)\right]$$

$$M_1^D = \tilde{\pi}_1$$

$$P_1^D = 0$$

We can verify that M_t^D is indeed a Martingale by induction (it is verified at $t = 1$)

$$\begin{aligned}
 \mathbb{E}_D[M_{t+1}^D | \mathcal{F}_t] &= \mathbb{E}_D[\tilde{\pi}_{t+1} - P_{t+1}^D | \mathcal{F}_t] \\
 &= \mathbb{E}_D[\tilde{\pi}_{t+1} | \mathcal{F}_t] - P_{t+1}^D \\
 &= \mathbb{E}_D[I_{t+1}\varphi_{AI}(\tilde{\pi}_t) + (1 - I_{t+1})\varphi_{AV}(\tilde{\pi}_t) | \mathcal{F}_t] - P_{t+1}^D \\
 &= \tilde{\pi}_t + \underbrace{(1 - \tilde{\pi}_t)(p_D - p_H)[\varphi_{AI}(\tilde{\pi}_t) - \varphi_{AV}(\tilde{\pi}_t)]}_{=-P_t^D} - P_{t+1}^D \\
 &= \tilde{\pi}_t - P_t^D \\
 &= M_t^D
 \end{aligned}$$

Finally, because of the unicity of the Doob-Meyer decomposition, the above decomposition is the only one.

The same reasoning yields the decomposition under state of the world H

$$\begin{aligned}
 \tilde{\pi}_t &= M_t^H + P_t^H \\
 M_t^H &\text{ is a martingale and } P_t^H \text{ is predictable and increasing} \\
 P_t^H &= (p_D - p_H) \sum_{i=1}^{t-1} (1 - \tilde{\pi}_i) [\varphi_{AI}(\tilde{\pi}_i) - \varphi_{AV}(\tilde{\pi}_i)] \\
 M_1^H &= \tilde{\pi}_1 \\
 P_1^H &= 0.
 \end{aligned}$$

Positivity of P_t^D and negativity of P_t^H are direct consequences of inequalities (3.2). □

3.6.1.13 Proof of Proposition 3.9

Proof. From the formulas for the updating functions

$$\varphi_{AI} = \frac{p_D \pi}{\bar{p}(\pi)} \quad \text{and} \quad \varphi_{AV} = \frac{(1 - p_D) \pi}{(1 - \bar{p}(\pi))},$$

we derive the derivatives with respect to p_D and p_H

$$\begin{aligned} \frac{\partial \varphi_{AI}}{\partial p_D} &= \frac{p_H \pi (1 - \pi)}{\bar{p}(\pi)^2} > 0 & \text{and} & \quad \frac{\partial \varphi_{AV}}{\partial p_D} = \frac{-(1 - p_H) \pi (1 - \pi)}{(1 - \bar{p}(\pi))^2} < 0, \\ \frac{\partial \varphi_{AI}}{\partial p_H} &= \frac{-p_D \pi (1 - \pi)}{\bar{p}(\pi)^2} < 0 & \text{and} & \quad \frac{\partial \varphi_{AV}}{\partial p_H} = \frac{(1 - p_D) \pi (1 - \pi)}{(1 - \bar{p}(\pi))^2} > 0. \end{aligned}$$

Finally, taking the derivatives of P_t^H and P_t^D and using the above derivatives yields the final result. □

Chapter 4

You'll Give Yourself a Bad Name: Reputation Effects in Repeated Audits ¹

¹This chapter a joint work with Pierre Picard. We are grateful to conference participants at EGRIE annual meetings for helpful discussions and comments.

Abstract

In principal-agent problems, the repetition of interactions in a dynamic setting may alter the equilibrium outcomes. In insurance fraud, the frequency of auditor-auditee interactions is higher when there is collusion between policyholders and service providers (e.g., car repairers, health care providers...). The same service provider usually handles claims filed by many policyholders affiliated to the same insurer, and thus the insurer-service provider interactions are repeated with reputation effects. We analyze this issue in a repeated game where the insurer may potentially face a dishonest service provider who colludes with policyholders. The insurer has beliefs about the type (honest or dishonest) of the service provider and she may verify the truthfulness of the claim through costly audits. The reputation of the service provider corresponds to these beliefs and changes over time, and misbehaving deteriorates this reputation. In the end, it may lead to a breach of contract and thus represents a threat that may deter from defrauding. We show that, at early periods, the insurer audits agents who would not be monitored in a static setting because their reputation is good enough. Corresponding dishonest agents who slipped under the radar and have an initially good reputation do not defraud systematically at early periods, as opposed to the instantaneous game. In addition, auditing efforts for medium reputations are lower as dishonest agents want to preserve the possibility of defrauding later. Both aspects corresponds to a reputation-based deterrence mechanism, where the fear of deteriorating one's reputation acts as a discipline device for dishonest service providers.

JEL Classification Numbers: C73, D82, D83, G22.

Keywords: Insurance fraud, Dynamic optimal auditing; Reputation; Learning; Deterrence.

4.1 Introduction

Many principal-agent relationships have a repeated nature, and from period to period, the principal may improve its knowledge of some hidden information of the agent. This situation arises in many problems of fraud inspection such as, for instance, the relationship between a firm and a subcontractor, or between a tax authority and a tax payer. In this paper, we consider the case of insurance fraud where insurers have repeated relationships with service providers who may be honest or dishonest.

Insurance fraud is a fundamental and complex problem for insurance markets. It is fundamental because it heavily distorts the risk-sharing mechanisms upon which insurance contracts rely, and without which insurance markets might collapse. It is also complex because auditing is motivated by mechanisms whose roles overlap and are hard to disentangle. Is an audit revealing no fraud synonymous of failure? Or does it just mean it completely deters fraud? How can information acquisition still motivate an audit if in the end deterrence is too strong to learn anything valuable?

The purpose of this paper is to combine both deterrence and informational aspects of auditing in a single framework to analyze their interactions. To do so, we model a repeated game between an insurer and a (honest or dishonest) auditee. The repeated character of interactions is particularly relevant when there is collusion between policyholders and service providers (SPs). As a single SP (e.g., car repairer or health service provider) may serve several policyholders of the same insurer, he interacts with the latter on a frequent basis. An SP acts as an expert who certifies² the claim, without which a fraud cannot occur, making an SP's complicity a necessary condition for cheating. Honest types may only certify invalid claims by mistake. Dishonest types on the other hand may defraud willingly by voluntarily certifying an illegitimate claim. Auditing is only profitable when fraudulent activity is large enough, and dishonest types rationally balance the risk of being caught with the potential gain from fraud. This is true even in a static setting when deterrence effects rely on the auditee's reaction to the threat of an audit. Informational effects come into play once interactions are repeated: based on the previous outcomes, the insurer has beliefs about the auditee's type, i.e., a reputation, on which the auditing decision depends.

Our main result shows that this combination exhibits a reputation-based deterrence effect that is two-fold: first, the set of auditees is larger at earlier periods and includes auditees whose reputation is good enough not to be audited in a static setting. Second, the insurer needs to

²i.e., that the policyholder has received a service and made a payment in return.

deploy less auditing efforts at earlier periods to deter auditees who would have been audited at later ones. Interestingly, this is the result of the interaction of two mechanisms, corresponding respectively to the deterrence and to the learning effects. With reputation effects, an auditee balances the intertemporal costs and benefits of cheating in the present: an audit today still threatens him as before, but being caught may also hinder the future opportunities of defrauding because the reputation worsens. The second mechanism of the reputation-based deterrence effects relies in the discriminatory power of auditing. Inspection allows the insurer to improve her information and to better audit in the future periods, or even exclude a dishonest auditee from his affiliation network.

These aspects relate to several strands of the academic literature. First and foremost, this article contributes to the literature on optimal auditing for insurance fraud.³ Following the work of [Townsend \(1979\)](#) and [Gale and Hellwig \(1985\)](#) on the Costly State Verification approach, it considers the problem of conducting an audit when the latter is costly. It tackles questions related to the relevance of random auditing, thus expanding the analysis of [Mookherjee and Png \(1989\)](#) by conditioning the levels of auditing on the reputation of auditees. Extending the idea of Costly State Falsification developed by [Crocker and Morgan \(1998\)](#), it exhibits a form of opportunity cost for the auditee to not defraud in order to mislead the auditor into believing he is not dishonest. More specifically, our work relates to the analysis of the different roles of auditing, besides recovering illegitimate payments. [Dionne, Giuliano and Picard \(2008\)](#) focus on auditing as a deterrence tool in a static Stackelberg game, and show that some claims should still be audited in spite of being individually unprofitable to audit. In [Aboutajdine and Picard \(2018\)](#) and [Aboutajdine and Picard \(2019b\)](#), we consider an alternative situation where there is no deterrence⁴ but the auditor can learn about the auditee's type in a multiperiod framework. The present article addresses the problem of combining both deterrence and information in a single framework.

Second, an important feature of our model, namely being a repeated game with monitoring, relates our work to the inspection games literature (see [Avenhaus, Von Stengel and Zamir \(2002\)](#) for a literature review). Originally applied to arms control and disarmament, these models have also been used for accounting and auditing problems in economics. [Borch \(1990\)](#) evokes the potential application of inspection games to the economics of insurance, specifically to ex-ante moral hazard problems where the insured has to take special measures to avoid accidents. In

³[Picard \(2013\)](#) offers a comprehensive review of this literature.

⁴i.e., the type dependent behavior of auditees is given and not affected by the risk of being spotted.

our setting, the problem is one of ex-post moral hazard (submitting a claim, valid or invalid, after the occurrence of a potential accident). Inspection games have also been applied to tax evasion settings. [Rubinstein \(1979\)](#) examines a case where an audit might unjustly target an honest auditee and shows that, while a one-shot game gives no choice but to set a large penalty, a repeated game allows for a more lenient policy that also induces the auditee to comply. [Greenberg \(1984\)](#) proposes an auditing scheme where individuals are separated into different groups with different auditing levels. Auditees move between groups depending on the previous cheating outcome. Another example of repeated game where cheating history is given by [Reinganum and Wilde \(1986\)](#). They apply the sequential equilibrium concept in a model where income reporting is a source of information for the tax inspector. More recently, [Varas, Marinovic and Skrzypacz \(2018\)](#) analyze a firm's reputation for quality problem, with an inspector who can conduct costly inspections. Notably, in their model, monitoring plays both an incentive and an informational role. Finally, the existence of an optimal information acquisition dimension to the insurer's problem relates it to the exploration/exploitation dilemma. It places our work in the continuity of the strategic experimentation literature and bandit problems (see [Bergemann and Välimäki \(2008\)](#)). In [Aboutajdine and Picard \(2019b\)](#), we show that the optimal auditing with learning but no deterrence problem is akin to a regular one-armed bandit problem where playing the risky arm corresponds to auditing, and playing the safe arm to not auditing. This is because the auditees are non-strategic and honest and dishonest auditees always submit a claim. In this paper, we relax both these assumptions. The underlying bandit problem becomes far more complex as it involves information collection even when no action is taken, which is related to the restless bandit literature (introduced by [Whittle \(1988\)](#)). It also involves strategic arms whose rewards depend on the player's strategy (e.g., [Braverman et al. \(2017\)](#)). Both topics have been rarely tackled in the academic literature, especially in the case of strategic armed bandits.

The rest of the paper is as follows: Section [4.2](#) presents the model, Section [4.3](#) characterizes our main results and Section [4.4](#) shows how our problem is related to bandit problems. Section [4.5](#) concludes. All proofs are in the Appendix.

4.2 Model

4.2.1 Notations and main assumptions

We consider the relationship between an insurer (she) and a service provider (SP, he) during T periods indexed by $t = 1, \dots, T$. The SP is in touch with policyholders and he certifies their

claims (i.e., he asserts that he delivered the service required by the event described in the claim, at the price reported by the policyholder). The SP may certify at most one claim per period to the insurer (certifying more claims would signal that he behaves in a fraudulent way). The SP may be honest or dishonest, i.e. of type H or D , respectively, and the insurer has beliefs about the SP 's type.

Claims are valid or invalid, and invalid claims may have been filed involuntarily (by error) or voluntarily (by fraud). A claim is involuntary invalid when the policyholder wrongly but honestly thought that the loss at the origin of the claim should be covered by the insurance policy, and the service provider does not observe any wrongdoing. Claims are voluntary invalid when the policyholder and the SP collude: the policyholder files a claim certified by the SP in order to trigger the payment of an insurance indemnity not included in the contractual coverage (e.g., falsifying a claim by inventing an event that did not occur, or misrepresenting the circumstances of this event). Honest SPs only certify invalid claims involuntarily, while dishonest SPs may certify invalid claims knowing that they are invalid. In short, insurance fraud takes the form of invalid claims that are willfully transmitted to the insurer by policyholders, with the complicity of dishonest service providers.

Let p_0 and p_H be the probability for an SP to certify a valid claim or to certify involuntarily an invalid claim, whatever the SP's type, with $p_0 + p_H < 1$. Let $y_t \in [0, 1]$ the probability for a dishonest SP to wilfully certify an invalid claim at period t when he has no claim to truthfully certify, which occurs with probability $1 - p_0 - p_H$. In what follows, y_t will be referred to as the fraud probability of a type D SP, which corresponds to the voluntary certification of an invalid claim with probability $y_t(1 - p_0 - p_H)$ when the SP is of type D . The insurer may audit a claim, which costs c to her, and we denote $x_t \in [0, 1]$ the probability of auditing a claim filed at period t . Audit reveals whether the claim is valid or not. Usually, in the case of an invalid claim, audit does not allow the insurer to establish whether this was intentional or unintentional, i.e., whether this was fraud or simply an error. In both cases, no indemnity is paid to the policyholder. When a fraudulent invalid claim is audited, the insurer may detect fraud unambiguously with probability $\alpha \in [0, 1]$.

Remark 1. *In a formally equivalent way, our model may also be interpreted as corresponding to the relationship between a firm and a sub-contractor supplying spare parts, assembled by the firm into a finished product. Each year, the firm expresses its needs under the form of new product specifications (e.g., additional safety or quality criteria), and the sub-contractor*

answers by transmitting a forecast report about expected additional costs induced by the change in product specifications. In this interpretation of the model, p_0 is the probability that the change in specifications actually induces unavoidable additional costs. However, with probability p_H the additional expected costs reported by the sub-contractor simply reflect a misunderstanding of the firm's requirements, or the fact that it is unaware of costless ways to comply with the firm's requirement (e.g., by using different inputs or by subcontracting itself with another firm for some parts of the product to be supplied). Such misunderstandings may be cleared by the firm through an audit of the sub-contractor's report. However, the announced additional costs may also result from the willingness of a dishonest sub-contractor to falsify its true cost, possibly through a cost-accounting manipulation, which may be detected with probability α if the forecast report is audited.

We wish to focus attention on the case where proving fraud is difficult, possibly impossible, and thus we postulate that α is small, possibly zero. More precisely, we assume

$$0 \leq \alpha < \frac{p_0}{p_0 + p_H}. \quad (4.1)$$

Each claim corresponds to a contractual insurance indemnity I , while defrauding costs ω to a dishonest SP. We may interpret ω as the expenses incurred by the SP and by the policyholder to falsify the claim (i.e., to prevent the insurer to immediately detect that the claim is invalid without any audit). Hence $I - \omega$ is the collusive gain in case of undetected fraud, to be shared between the dishonest SP and the policyholder. For simplicity, we assume that collusion takes the form of a take-or-leave-it offer made by the SP to the policyholder, with the whole collusive surplus $I - \omega$ being allocated to the SP, but this is an unimportant assumption in our analysis. More importantly, we assume:

$$c < (1 - p_0)I, \quad (4.2)$$

$$(p_0 + p_H)c > p_H I. \quad (4.3)$$

According to assumption (4.2), if the insurer knows with certainty that the SP is of type D and that he systematically defrauds (i.e., $y = 1$), then it is profitable to audit the claim, because the audit cost c is lower than the expected recouped indemnity $(1 - p_0)I$. Conversely, according to assumption (4.3), if the insurer knows with certainty that the SP is of type H , then it is not worthwhile to audit the claim, because the audit cost c is larger than the expected recouped

indemnity $p_H I / (p_0 + p_H)$ associated with involuntary errors.

The insurer offers a multi-period contract to the SP at the beginning of period 1, whereby the SP is allowed to act as an insurance intermediary from $t = 1$ to $t = T$, which consists in certifying claims filed by policyholders. At period 1, the insurer has initial beliefs $\pi_1 \in (0, 1)$ about the SP's type: π_1 is the probability that the SP is of type D , and thus $1 - \pi_1$ is the probability that he is of type H .⁵ At each period $t = 2, \dots, T$, the insurer has updated beliefs π_t deduced from π_{t-1} and from the scenario that occurred at period $t - 1$. A breach of contract takes place at the beginning of period t only if the insurer can prove unambiguously that the SP is dishonest.⁶ In such a case, another contract is offered to a new SP for the remaining periods $t, t + 1, \dots, T$, with insurer's beliefs $\bar{\pi} \in (0, 1)$ about this new SP.

For any fraudulent claim filed at period t , the expected profit of the dishonest SP involved in the collusion is the difference between his gain $I - \omega$ if there is no audit, and his loss ω in the case of an audit, weighted by the probabilities $1 - x_t$ and x_t , respectively. Since a type D SP defrauds with probability $y_t(1 - p_0 - p_H)$. We may denote

$$\Pi(x_t, y_t) = y_t(1 - p_0 - p_H)[(1 - x_t)(I - \omega) - x_t\omega], \quad (4.4)$$

his expected profit drawn from fraud at period t .

We also denote

$$C(x_t, y_t, \pi_t) = I[p_0 + K(\pi_t y_t)] - x_t[IK(\pi_t y_t) - c(p_0 + K(\pi_t y_t))] \quad (4.5)$$

the expected cost of the insurer at period t , where

$$K(\pi_t y_t) = p_H + \pi_t y_t(1 - p_0 - p_H) \quad (4.6)$$

is the probability of an invalid claim at period t when the SP is of type D with probability π_t . This probability is the sum of p_H , which is the probability of an invalid claim transmitted by error, and of $\pi_t y_t(1 - p_0 - p_H)$ which is the probability of a fraudulent claim. The expected insurer's cost $C(x_t, y_t, \pi_t)$ is the difference between the expected insurance indemnity $I[p_0 + K(\pi_t y_t)]$ and the expected net proceeds from audit $x_t[IK(\pi_t y_t) - c(p_0 + K(\pi_t y_t))]$.

⁵This may correspond to the fraction of dishonest individuals in a population, in which the SP is randomly drawn. Initial beliefs may also be affected by business references or by other signals informative about the SP's honesty.

⁶A breach of contract without evidence of dishonesty would lead to bad faith penalty imposed by a court.

Five scenarios may occur at each each period $t = 1, \dots, T$, denoted $s_t \in S$ with

$$S = \{AF, AI, AV, NA, NC\}.$$

These scenarios are detailed in Table 4.1. The objective of the insurer is to minimize her discounted expected cost, and the objective of a type D SP is to maximize his discounted expected profit from fraud. We denote $\delta \in (0, 1]$ the discounting factor, common to the insurer and to the SP.

Table 4.1 – POSSIBLE SCENARIOS AT A GIVEN PERIOD

Scenario	Details	Auditor's Cost
AF	A claim is filed, audit reveals an invalid claim, fraud is proved	c
AI	A claim is filed, audit reveals an invalid claim, no fraud is proved	c
AV	A claim is filed, audit reveals a valid claim	$I + c$
NA	A claim is filed, no audit is performed	I
NC	No claim is filed	0

4.2.2 Definition of an equilibrium

Let $h_t = (s_1, s_2, \dots, s_{t-1}) \in S^t$ be the history of the relationship between the insurer and the SP before period t , with $t \geq 1$. This history affects the expected profit of a type D SP and the expected cost of the insurer only through its effect on beliefs π_t . We make the following assumptions about strategies and beliefs:

1. Strategies are Markovian: they define audit and fraud decisions at period t as functions of the insurer's beliefs at the same period, and they are denoted $x_t(\pi_t)$ and $y_t(\pi_t)$ for all $t = 1, \dots, T$.
2. Beliefs π_t depend on previous beliefs π_{t-1} and on the scenario s_{t-1} that occurred at the previous period, and we will write them as: $\pi_t = \pi_t(\pi_{t-1}, s_{t-1})$ for all $t = 2, \dots, T$.

Definition 4.1. *A Markov Subgame Perfect Bayesian Equilibrium (in short, an Equilibrium) is characterized by strategies $x_t^*(\cdot) : [0, 1] \rightarrow [0, 1]$ and $y_t^*(\cdot) : [0, 1] \rightarrow [0, 1]$ for all $t = 1, \dots, T$ defining actions $x_t^*(\pi_t), y_t^*(\pi_t)$ as functions of beliefs, and by belief updating functions $\pi_t^*(\cdot) : [0, 1] \times S \rightarrow [0, 1]$ for all $t = 2, \dots, T$ defining updated beliefs $\pi_t^*(\pi_{t-1}, s_{t-1})$ as functions of previous beliefs and previous period scenario, such that:*

1. For all $t = 1, \dots, T$ and all belief $\pi_t \in [0, 1]$, $x_t^*(\pi_t)$ minimizes

$$C_t(x_t, y_t^*(\pi_t), \pi_t) = C(x_t, y_t^*(\pi_t), \pi_t) + \sum_{\theta=1}^{T-t} \delta^\theta \mathbb{E}_{\tilde{\pi}_{t+\theta}} \left[C(x_{t+\theta}^*(\tilde{\pi}_{t+\theta}), y_{t+\theta}^*(\tilde{\pi}_{t+\theta}), \tilde{\pi}_{t+\theta}) \right],$$

w.r.t. $x_t \in [0, 1]$, where, for all $\theta = 1, \dots, T - t$, $\tilde{\pi}_{t+\theta}$ denotes the random beliefs deduced from strategies $x_t, x_{t+1}^*(\cdot), \dots, x_T^*(\cdot)$ and $y_t^*(\cdot), y_{t+1}^*(\cdot), \dots, y_T^*(\cdot)$, and from the belief updating functions $\pi_{t+1}^*(\cdot), \dots, \pi_T^*(\cdot)$.⁷

2. For all $t = 1, \dots, T$ and all beliefs $\pi_t \in [0, 1]$, $y_t^*(\pi_t)$ maximizes

$$\Pi_t(x_t^*(\pi_t), y_t, \pi_t) = \Pi(x_t^*(\pi_t), y_t, \pi_t) + \sum_{\theta=1}^{T-t} \delta^\theta \mathbb{E}_{\tilde{\pi}_{t+\theta}} \left[\Pi(x_{t+\theta}^*(\tilde{\pi}_{t+\theta}), y_{t+\theta}^*(\tilde{\pi}_{t+\theta}), \tilde{\pi}_{t+\theta}) \right],$$

w.r.t. $y_t \in [0, 1]$, where, for all $\theta = 1, \dots, T - t$, $\tilde{\pi}_{t+\theta}$ denote the random beliefs deduced from strategies $x_t^*(\cdot), x_{t+1}^*(\cdot), \dots, x_T^*(\cdot)$ and $y_t, y_{t+1}^*(\cdot), \dots, y_T^*(\cdot)$, and from the belief updating functions $\pi_{t+1}^*(\cdot), \dots, \pi_T^*(\cdot)$.

3. For all $t = 1, \dots, T$ and all (π_{t-1}, s_{t-1}) that can be reached with positive probability on the equilibrium path, the belief updating functions $\pi_t^*(\pi_{t-1}, s_{t-1})$ are deduced from strategies $x_{t-1}^*(\pi_{t-1}), y_{t-1}^*(\pi_{t-1})$ through Bayes Law when $s_{t-1} \neq AF$, and $\pi_t^*(\pi_{t-1}, s_{t-1}) = \bar{\pi}$ when $s_{t-1} = AF$.

Hereafter, for all $t = 1, \dots, T$, the optimized cost and payoff are defined as functions of π_t alone:

$$C_t^*(\pi_t) = C_t(x_t^*(\pi_t), y_t^*(\pi_t), \pi_t) \quad \text{and} \quad \Pi_t^*(\pi_t) = \Pi_t(x_t^*(\pi_t), y_t^*(\pi_t), \pi_t).$$

For simplicity of comparison, we can also define their average discounted counterparts as

$$\bar{C}_t^*(\pi_t) = \frac{1}{\sum_{\theta=0}^{T-t} \delta^\theta} C_t^*(\pi_t) \quad \text{and} \quad \bar{\Pi}_t^*(\pi_t) = \frac{1}{\sum_{\theta=0}^{T-t} \delta^\theta} \Pi_t^*(\pi_t).$$

4.2.3 Instantaneous game

Let us start with the one-period model, i.e., $T = 1$ without repeated interaction between insurer and SP. We refer to this case as the *instantaneous game*. Its equilibrium is characterized in Proposition 4.1.

Proposition 4.1. *An equilibrium of the instantaneous game (i.e., when $T = 1$) is given, for an*

⁷In other words, $\tilde{\pi}_{t+\theta} = \pi_{t+\theta}^*(\tilde{\pi}_{t+\theta-1}, \tilde{s}_{t+\theta-1})$ for $\theta = 1, \dots, T - t$, where $\tilde{s}_{t+\theta-1}$ is distributed in S with probabilities induced by strategies $x_{t+\theta-1}^*(\tilde{\pi}_{t+\theta}), y_{t+\theta-1}^*(\tilde{\pi}_{t+\theta})$ when $\theta = 2, \dots, T - 1$ and by strategies $x_t, y_t^*(\pi_t)$ when $\theta = 1$. Similarly for $\tilde{\pi}_{t+\theta}$.

initial belief π at the beginning of the period, by $x_1^*(\pi) = \hat{x}(\pi)$, $y_1^*(\pi) = \hat{y}(\pi)$, where

$$\hat{x}(\pi) = \hat{x} \text{ and } \hat{y}(\pi) = \frac{\hat{\pi}}{\pi} \text{ if } \pi > \hat{\pi},$$

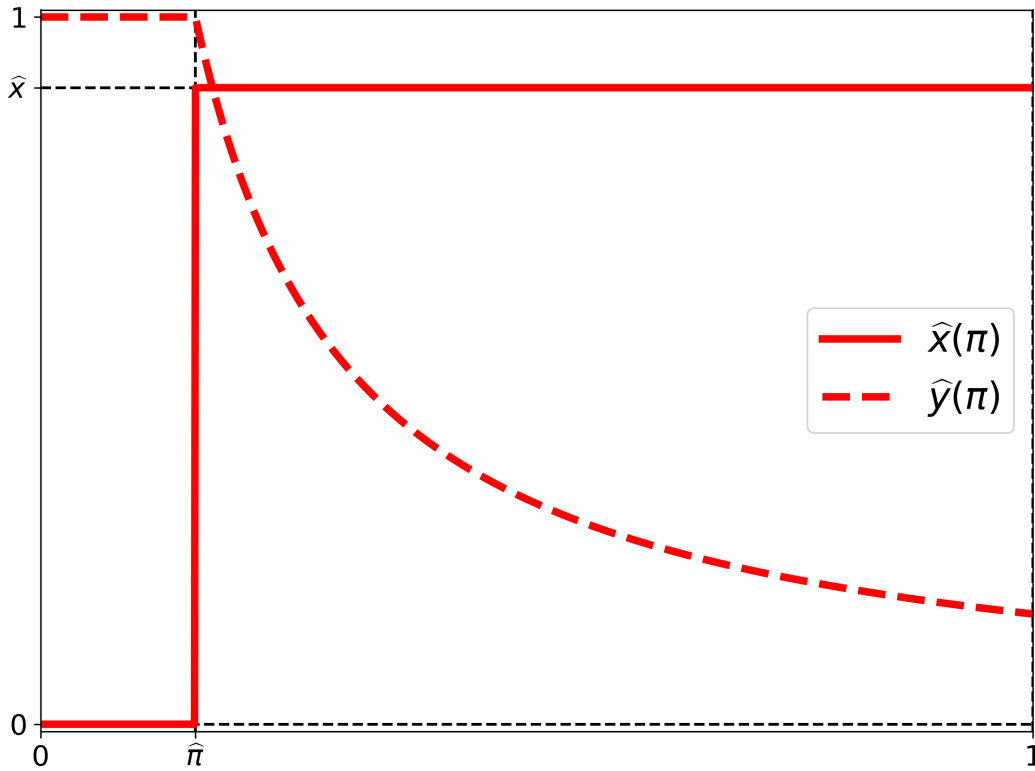
$$\hat{x}(\pi) = 0 \text{ and } \hat{y}(\pi) = 1 \text{ if } \pi \leq \hat{\pi},$$

where

$$\hat{\pi} = \frac{(p_0 + p_H)c - p_H I}{(1 - p_0 - p_H)(I - c)} \in (0, 1), \quad (4.7)$$

$$\hat{x} = \frac{I - \omega}{I} \in (0, 1). \quad (4.8)$$

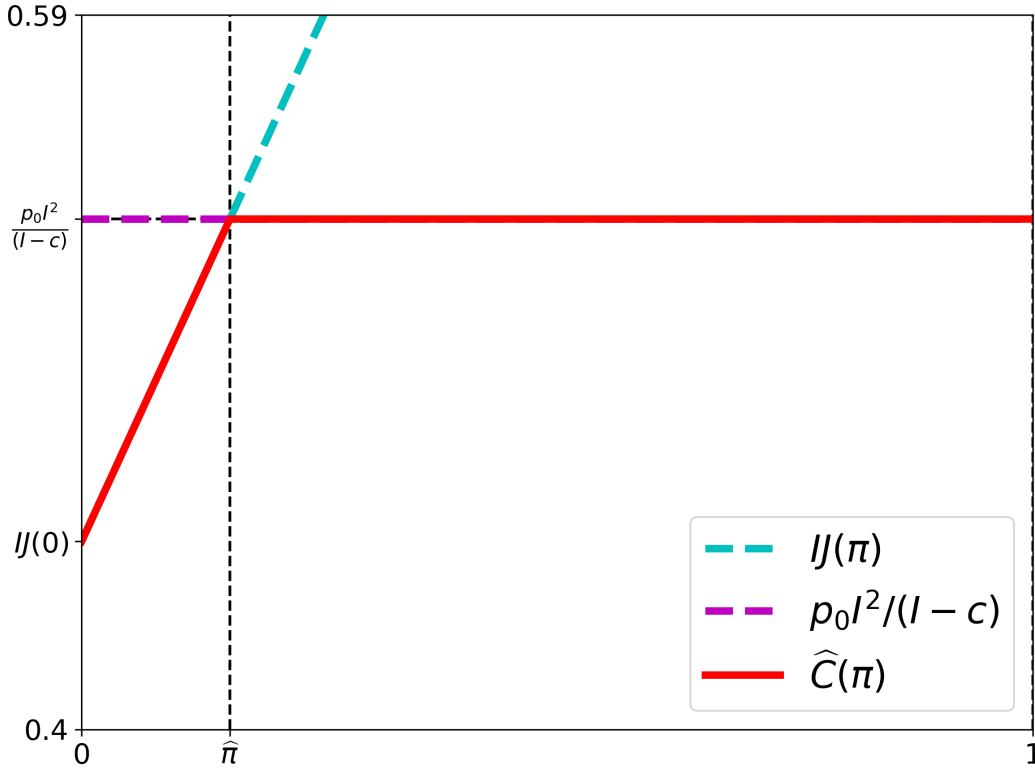
Figure 4.1 – INSTANTANEOUS EQUILIBRIUM AUDITING AND FRAUD



Proposition 4.1 is very intuitive. $\hat{\pi}$ is the threshold such that, when $\pi < \hat{\pi}$, there is not enough incentives for triggering an audit, and a type D SP always defraud, i.e., $x_1^*(\pi) = 0$ and $y_1^*(\pi) = 1$. When $\pi > \hat{\pi}$, a type D SP should defraud with probability less than one, for otherwise the insurer would systematically trigger an audit. In that case, the fraud probability $y_1^*(\pi) = \hat{y}(\pi) = \hat{\pi}/\pi \in (0, 1)$ makes the insurer indifferent between auditing and not auditing. Symmetrically, the equilibrium audit probability $x_1^*(\pi) = \hat{x} \in (0, 1)$ makes the type D SP indifferent defrauding and

not defrauding. The equilibrium is thus in mixed strategies when $\pi > \hat{\pi}$.

Figure 4.2 – INSTANTANEOUS EQUILIBRIUM COST



A simple calculation shows that the equilibrium cost of the insurer may be written as

$$\hat{C}(\pi) = C(\hat{x}(\pi), \hat{y}(\pi), \pi) = \begin{cases} IJ(\pi) & \text{if } \pi < \hat{\pi}, \\ p_0I^2/(I - c) & \text{if } \pi \geq \hat{\pi}, \end{cases} \quad (4.9)$$

where

$$J(\pi) = p_0 + p_H + \pi(1 - p_0 - p_H)$$

or equivalently

$$\hat{C}(\pi) = C(\hat{x}(\pi), \hat{y}(\pi), \pi) = \inf\{IJ(\pi), p_0I^2/(I - c)\}. \quad (4.10)$$

Hence, when $\pi < \hat{\pi}$, there is no audit and the expected costs increases linearly with the probability of a type D SP. When $\pi \geq \hat{\pi}$, the reputation of the SP is bad enough for auditing with positive probability to be worthwhile, and the equilibrium expected cost reaches an upper bound $p_0I^2/(I - c)$.

4.3 Reputation-based Deterrence

In the following we are interested in multiperiod settings where $T \geq 2$. In Section 4.3.1, we explain how the insurer refines his beliefs through Bayesian updating and how this may induce players to deviate from the myopic strategies. We also define two types of reputation-based deterrence effects. In Section 4.3.2, we solve the game for two periods and show that only one type of reputation-based deterrence occurs, unless fraud can be proved unambiguously. Section 4.3 shows that, with three periods, both types of reputation-based deterrence take place, regardless of the possibility to unambiguously prove fraud. Finally, Section 4.3.4 gives sufficient conditions for the first type of reputation-based deterrence effect to happen for more than three periods.

4.3.1 Learning and the inter-temporal threat of an audit

At the end of period t , whether a claim was submitted or not and the outcome of a potential audit allow the insurer to refine his beliefs about the SP's type.

If a claim was submitted and an audit reveals unambiguously that the claim was fraudulent, the contract between the SP and the insurer is breached. In that case, another contract is offered for period $t + 1$ to a new SP. Since the insurer has beliefs $\bar{\pi}$ about the type of this new SP, we have

$$\pi_{t+1} = \bar{\pi} \text{ if } s_t = AF.$$

Using Bayes law yields the updated beliefs for all other scenarios for initial beliefs π_t and strategy $y_t \equiv y_t(\pi_t)$. Simple calculations yield

$$\pi_{t+1} = \varphi_{AF}(\pi_t, y_t) = \bar{\pi} \quad \text{when } s_t = AF,$$

$$\pi_{t+1} = \varphi_{AV}(\pi_t, y_t) = \pi_t \quad \text{when } s_t = AV,$$

$$\pi_{t+1} = \varphi_{AI}(\pi_t, y_t) = \frac{\pi_t[p_H + y_t(1 - \alpha)(1 - p_0 - p_H)]}{p_H + \pi_t y_t(1 - \alpha)(1 - p_0 - p_H)} \quad \text{when } s_t = AI,$$

$$\pi_{t+1} = \varphi_{NA}(\pi_t, y_t) = \frac{\pi_t[p_0 + p_H + y_t(1 - p_0 - p_H)]}{p_0 + p_H + \pi_t y_t(1 - p_0 - p_H)} \quad \text{when } s_t = NA,$$

$$\pi_{t+1} = \varphi_{NC}(\pi_t, y_t) = \frac{\pi_t(1 - y_t)}{1 - \pi_t y_t} \quad \text{when } s_t = NC.$$

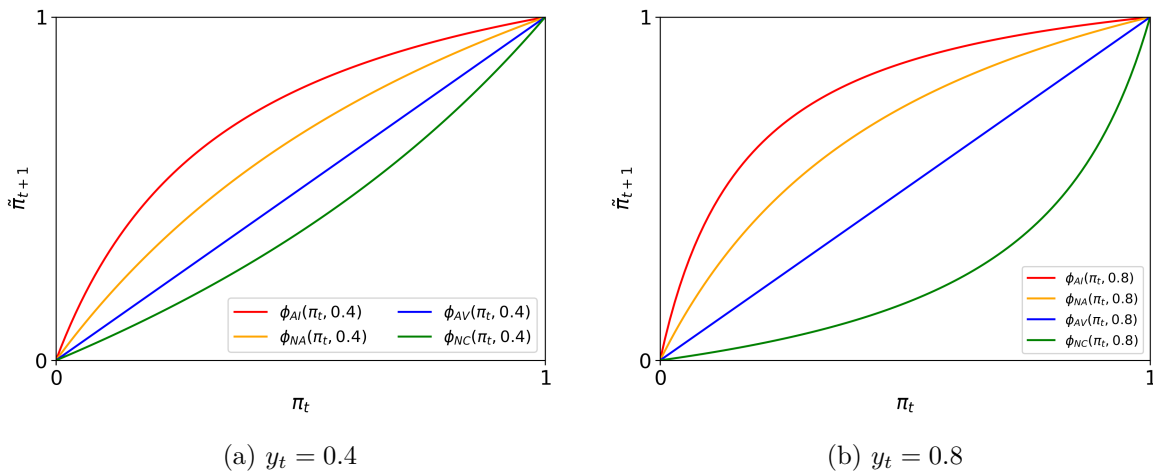
Functions $\varphi_{AI}(\pi_t, y_t)$, $\varphi_{NA}(\pi_t, y_t)$ are increasing and concave w.r.t. π_t , and $\varphi_{NC}(\pi_t, y_t)$ is increasing and convex w.r.t. π_t , with

$$\varphi_{NC}(\pi_t, y_t) < \pi_t < \varphi_{NA}(\pi_t, y_t) < \varphi_{AI}(\pi_t, y_t), \quad (4.11)$$

for all $\pi_t \in (0, 1]$, $y_t \in [0, 1]$, where the last inequality results from assumption (4.2).

Figure 4.3 plots the updating functions for low fraud (Figure 4.3a, $y_t = 0.4$) and high fraud (Figure 4.3b, $y_t = 0.8$). The higher the fraud probability, the more different the updating functions, and the more information obtained from updating. It also shows how an SP can influence his reputation. By defrauding more, his reputation is damaged more when caught defrauding. In other words, there is a trade-off for the SP between immediate gains from fraud, and long-term gains from preserving his reputation.

Figure 4.3 – UPDATING FUNCTIONS FOR ($p_0 = 0.2, p_H = 0.1$)



Definition 4.2 (Reputation-based fraud deterrence). *In a multiperiod setting (i.e., $T \geq 2$), an equilibrium may exhibit a reputation-based fraud deterrence effect in the extensive and/or in the intensive margins.*

An equilibrium at period t exhibits a reputation-based fraud deterrence effect in the extensive margin when there exists an interval $(a, b]$, with $a < b \leq \hat{\pi}$, such that $x_t^(\pi_t) > 0$ and/or $y_t^*(\pi_t) < 1$ when $\pi_t \in (a, b]$.*

An equilibrium at period t exhibits a reputation-based fraud deterrence effect in the intensive margin when there exists an interval $[d, e)$, with $\hat{\pi} \leq d < e$, such that $x_t^(\pi_t) < \hat{x}$ and/or $y_t^*(\pi_t) < \hat{y}(\pi_t)$ when $\pi_t \in [d, e)$.*

In words, an equilibrium displays a reputation-based fraud deterrence effect in the extensive

margin when, at period $t \in \{1, \dots, T - 1\}$, the insurer audits claims with positive probability and/or the type D SP defrauds with probability less than one, under beliefs $\pi_t \in (0, \hat{\pi})$ for which there would be systematic fraud without audit at the equilibrium of the instantaneous game. In this case, auditing increases in the extensive margin as SPs with priors below the instantaneous profitability threshold $\hat{\pi}$ enter the auditing pool. Alternatively, fraud decreases in the extensive margin as SPs with priors below the instantaneous threshold $\hat{\pi}$ exit the systematic fraud pool.

As can be expected from the characterization of the instantaneous game equilibrium, when $\pi_t < \hat{\pi}$, auditing claims cannot be profitable to the insurer if her calculation is only based on the auditing cost at the same period. However, auditing also allows the insurer to improve her information about the SP's type, which may be useful to her in the next period. The threat of being more precisely identified as a type D at the beginning of the following period reduces the propensity of a dishonest SP to defraud, hence the possibility of an equilibrium where there is auditing and no-systematic fraud because of this learning process.

An equilibrium displays a reputation-based fraud deterrence effect in the intensive margin when, at period $t \in \{1, \dots, T - 1\}$, the insurer audits claims with positive but lower probability than the instantaneous auditing probability, and/or the type D SP defrauds with a probability lower than the instantaneous defrauding probability. In this case, auditing decreases in the intensive margin as the insurer needs to put in less auditing efforts to make a type D indifferent. Alternatively, fraud decreases in the intensive margin as a type D needs to defraud less to make the auditor indifferent.

Since optimal auditing and optimal fraud depend on the reputation π_t , the evolution of beliefs through learning may change the optimal inter-temporal auditing and defrauding strategies. Hereafter, we will primarily focus on how the set of beliefs π_t with a pure equilibrium $(x^*(\pi_t), y^*(\pi_t)) = (0, 1)$ evolves in time, i.e., on the first type of reputation-based deterrence. In that perspective, let us define the optimal thresholds $(\pi_t^*)_{t \in \{1, \dots, T\}}$ as follows.

Definition 4.3. *For any $t \in \{1, \dots, T\}$, let*

$$\pi_t^* = \sup \{ \pi \in [0, 1] \mid \forall \mu \in (0, \pi), (x_t^*(\mu), y_t^*(\mu)) = (0, 1) \} \quad (4.12)$$

π_t^* is the belief threshold under which there is no deterrence, i.e., the insurer does not audit and type D SPs always defraud. In particular, at time T , $\pi_T^* = \hat{\pi}$. This is a direct consequence of the fact that at the last period T , there is no subsequent period through which learning in T might influence optimal policies, and the equilibrium is the instantaneous one. In addition, the

existence of reputation-based deterrence in the extensive margin at time t is equivalent to having $\pi_t^* < \hat{\pi}$.

4.3.2 $T = 2$: Reputation-based deterrence in the intensive margin only

To establish the possibility of such a learning-based deterrence effect, we will first consider a two-period model, i.e., the case where $T = 2$. Period 2 is the last period, and thus the equilibrium of the instantaneous game described in Proposition 4.1 is also an equilibrium of at period 2. In other words, period 2 equilibrium strategies are

$$\begin{aligned} x_2^*(\pi_2) &= \hat{x}(\pi_2), \\ y_2^*(\pi_2) &= \hat{y}(\pi_2), \end{aligned}$$

where functions $\hat{x}(\cdot)$ and $\hat{y}(\cdot)$ correspond to the definition given in Proposition 4.1.

In addition to the current period myopic cost, the insurer takes into account the future cost that depends on the outcome of the current period, and on the corresponding updating transition. Let $y_1^*(\pi_1)$ be the period 1 equilibrium fraud and define the period 1 thresholds π_1^{NA} and π_1^{NC} as

$$\varphi_{NA}(\pi_1^{NA}, y_1^*(\pi_1^{NA})) = \hat{\pi} \quad \varphi_{NC}(\pi_1^{NC}, y_1^*(\pi_1^{NC})) = \hat{\pi}.$$

We have

$$\pi_1^{NA} < \hat{\pi} < \pi_1^{NC}.$$

4.3.2.1 Case where unambiguously proving fraud is impossible ($\alpha = 0$)

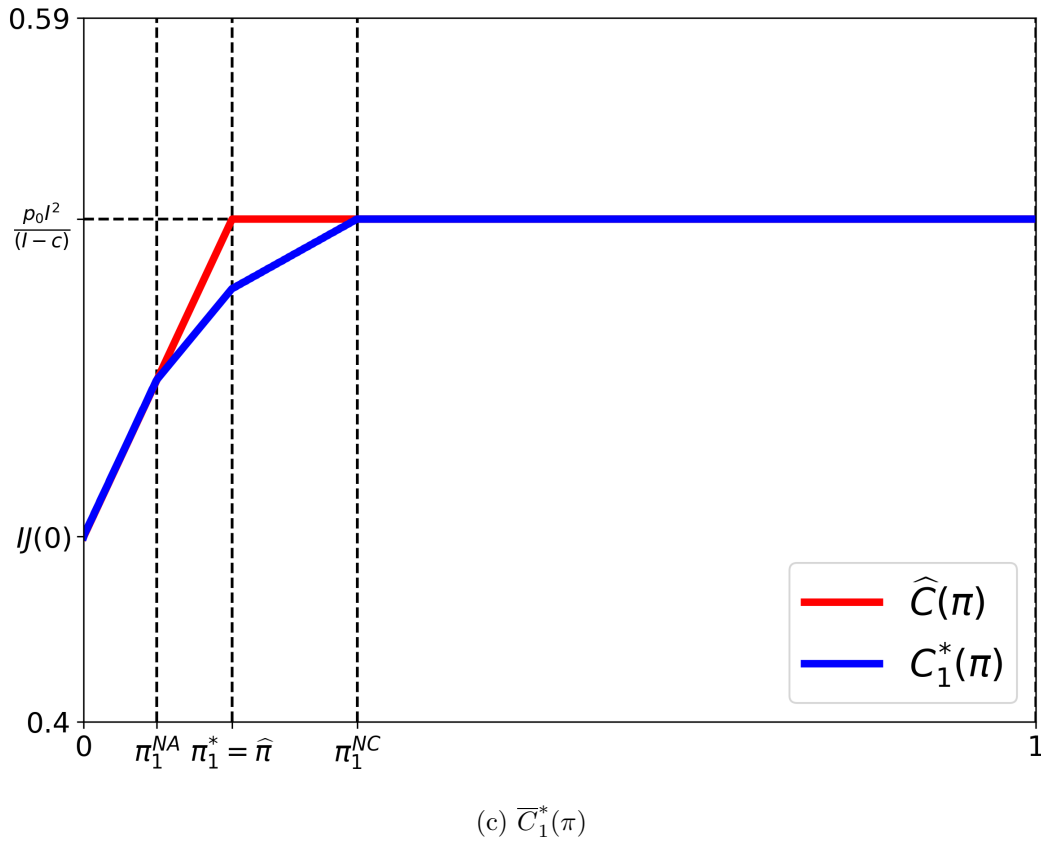
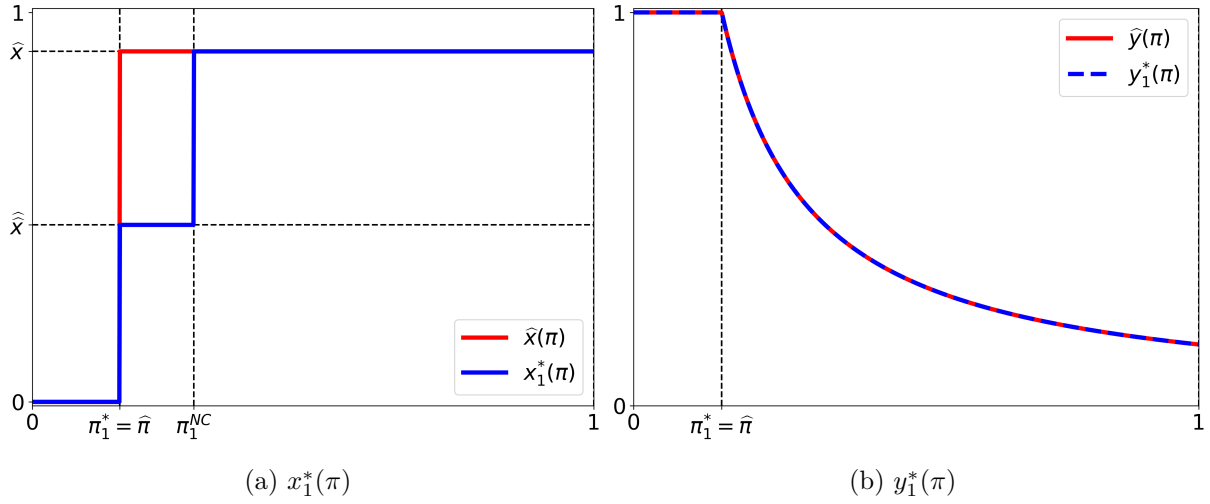
Proposition 4.2. *When $T = 2$ and $\alpha = 0$, the $t = 1$ equilibrium exhibits no reputation-based fraud deterrence effect in the extensive margin as $\pi_1^* = \hat{\pi}$. In addition, for all priors, equilibrium fraud is the instantaneous fraud $y_1^* = \hat{y}(\pi_1)$. However, the auditor makes less auditing efforts for intermediate reputations $\pi_1 \in (\hat{\pi}, \pi_1^{NC})$, as $x_1^*(\pi_1) = \hat{\hat{x}}$ on this interval, with*

$$\hat{\hat{x}} = \frac{(I - \omega)(1 - \delta(1 - p_0 - p_H))}{I} = (1 - \delta(1 - p_0 - p_H))\hat{x} < \hat{x}.$$

Thus, we have reputation-based fraud deterrence effects in the intensive margin.

Figures 4.4 represent the optimal audit, fraud and average discounted cost as functions of

Figure 4.4 – AUDIT, FRAUD AND AVERAGE DISCOUNTED COST AT $(t, T) = (1, 2)$



initial belief $\pi = \pi_1$. Figure 4.4a shows that there is no reputation-based deterrence in the extensive margin in the penultimate period as $\pi_{T-1}^* = \pi_1^* = \hat{\pi}$. On the opposite, reputation-based fraud deterrence in the intensive margin occurs, as $[d, e) = [\hat{\pi}, \pi_1^{NC})$. Indeed, the payoff of a type D has two components, a current period one, equivalent to the myopic payoff, and a subsequent period one, equal to $\delta(1 - y_1)(1 - p_0 - p_H)$. The derivative of the current period component is equal to 0 for $x_1 = \hat{x}$ from Proposition 4.1. Thus the total derivative is equal to $-\delta(1 - p_0 - p_H)$, independent of x_1 . Consequently, for $x_1 = \hat{x}$, type D 's best response is $y_1 = 0$, and the myopic

equilibrium cannot be an equilibrium at $t = 1$ in $(\hat{\pi}, \pi_1^{NC})$. Figure 4.4b illustrates the fact that the optimal fraud level at equilibrium does not change. Figure 4.4c illustrates the equilibrium optimal average discounted cost and shows the decrease in cost thanks to the learning effect. Intuitively, this decrease occurs for intermediate reputations (π_1^{NA}, π_1^{NC}) for which uncertainty reduction is the most effective.

4.3.2.2 Case where it may be possible to unambiguously prove fraud ($\alpha > 0$)

Before considering a $T = 3$ setting, we focus on the case where fraud can be proven unambiguously, i.e., $\alpha > 0$.

Proposition 4.3. *When $T=2$, the equilibrium exhibits a reputation-based fraud deterrence effect in the extensive margin when $\alpha > 0$ and $\bar{\pi} < \hat{\pi}$.*

In other words, the proof of Proposition 4.3 shows that there is an interval $(\pi_1^\alpha, \hat{\pi})$ where $y_1^*(\pi) < 1$ and $x_1^*(\pi) > 0$, meaning auditing to trigger deterrence effects is an equilibrium strategy, while there would have been no equilibrium auditing in this same interval in a purely instantaneous setting. This is a very intuitive result: when there is a positive probability to unambiguously detect fraud, the possibility to replace the defrauder by another SP with better reputation provides an additional incentive to audit claims at period 1. This creates a credible threat for type D SPs, and consequently fraud is reduced and auditing occurs with positive probability for some π_1 smaller than $\hat{\pi}$.

4.3.3 $T = 3$: Reputation-based deterrence in the extensive and the intensive margins

As for Section 4.3.2, it will be useful to define some thresholds. First, because of the notation, the first period from the $T = 2$ case becomes the second one, and the previous thresholds are now denoted π_2^{NA} and π_2^{NC} . Let $y_1^*(\pi_1)$ be the period 1 equilibrium fraud when $T = 3$ and define the period 1 thresholds π_1^{NA} and π_1^{NC} as

$$\varphi_{NA}(\pi_1^{NA}, y_1^*(\pi_1^{NA})) = \pi_2^{NA} \quad \varphi_{NC}(\pi_1^{NC}, y_1^*(\pi_1^{NC})) = \pi_1^{NC}.$$

We have

$$\pi_1^{NA} < \pi_2^{NA} < \hat{\pi} < \pi_2^{NC} < \pi_1^{NC}.$$

Proposition 4.4. *Assume $\bar{\pi} < \hat{\pi}$, and*

$$p_0 + p_H < \frac{1 - \hat{\pi}}{2 - \hat{\pi}}, \quad (4.13)$$

$$1 - \sqrt{\frac{c}{I}} < \frac{I - c}{2I - c}. \quad (4.14)$$

Then, for $T = 3$, the $t = 1$ equilibrium exhibits a reputation-based fraud deterrence effect in the extensive margin regardless of α as $\pi_1^ < \hat{\pi}$.*

Condition (4.13) is a sufficient condition for the proposition to hold. It is made to make the proof easier, but it is not necessary. It may be interpreted as an assumption about the importance of fraud: the probability that a type D is in a position to defraud is $1 - p_0 - p_H$, and according to (4.13), this probability should be larger than $1/(2 - \hat{\pi})$. (4.14) is a condition on I and c under which (4.2), (4.3) and (4.13) are compatible. Contrary to Proposition 4.3, it is valid also when $\alpha = 0$.

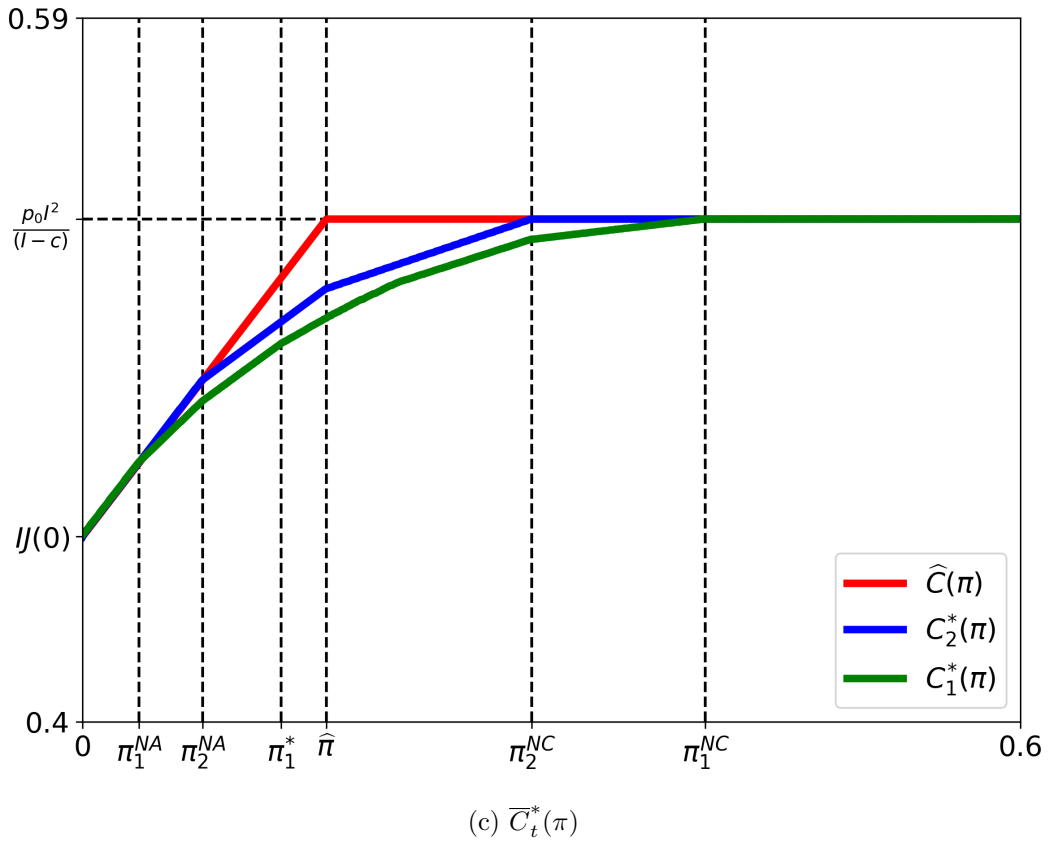
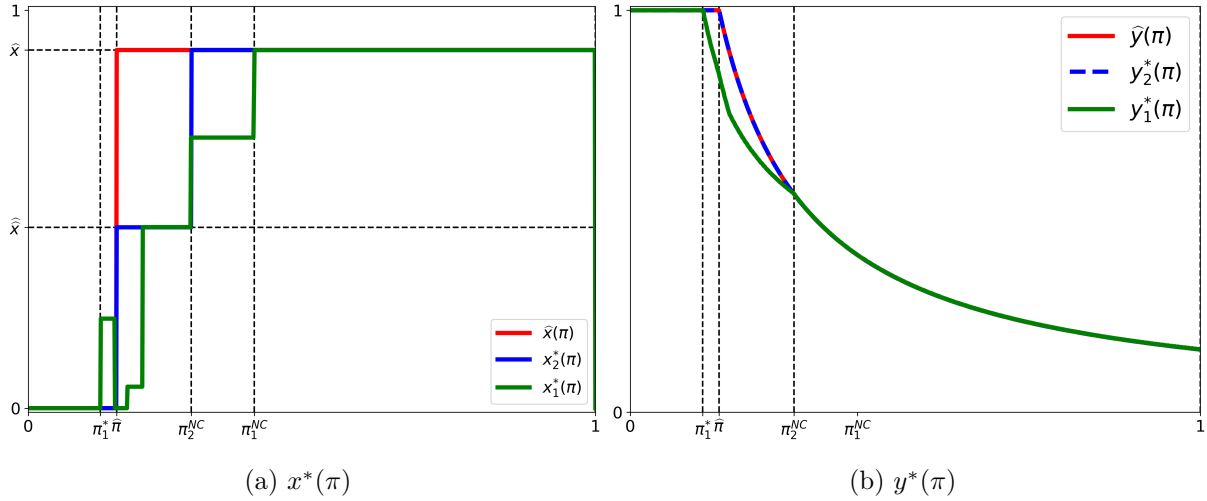
Figures 4.5 represent simulation⁸ results for the optimal audit, fraud and average discounted cost at the first period when there are three periods in total. Figure 4.5a shows that there is some reputation-based deterrence in the first period as $\pi_{T-2}^* = \pi_1^* < \hat{\pi}$. Figure 4.5b shows that the optimal fraud level decreases in comparison to the instantaneous benchmark on the interval (π_1^*, π_2^{NC}) as a result of this reputation-based deterrence effect. Figure 4.5c illustrates the equilibrium optimal average discounted cost and shows the decrease in cost thanks to the accumulation of learning in time. Notably, the set of beliefs for which a decrease occurs is larger in the first period (green curve) than in the second/penultimate period.

Therefore, both types of reputation-based deterrence effects come into play. First, for priors in the interval $[\hat{\pi}, \pi_1^{NC})$, auditing is lower than in the myopic equilibrium. More specifically, in the interval $[\hat{\pi}, \pi_2^{NC})$, fraud is also lower than in the myopic equilibrium. Both situations correspond to reputation-based deterrence in the intensive margin. Strikingly, the simulation results indicate that the mixed equilibrium auditing may drop to 0 above $\hat{\pi}$. Second, for priors in the interval $(\pi_{t+1}^*, \hat{\pi}]$, there is a mixed equilibrium with some fraud and non systematic deterrence. This is a reputation-based deterrence effect in the extensive margin, whereby priors $(\pi_{t+1}^*, \hat{\pi}]$ are included in the auditing set while they are not in the myopic equilibrium.

The proof relies on a very intuitive approach, where we look, for a prior $\pi_{T-2}^0 < \hat{\pi}$ and close enough to $\hat{\pi}$, at a difference in costs $\xi(\pi_{T-2}^0)$. This quantity is the difference between the

⁸For $(p_0, p_H, c, I) = (0.3, 0.15, 0.44, 1)$. Assumptions (4.2), (4.3), (4.13) and (4.14) are compatible for these values.

Figure 4.5 – AUDIT, FRAUD AND AVERAGE DISCOUNTED COST AT $(t, T) = (1, 3)$



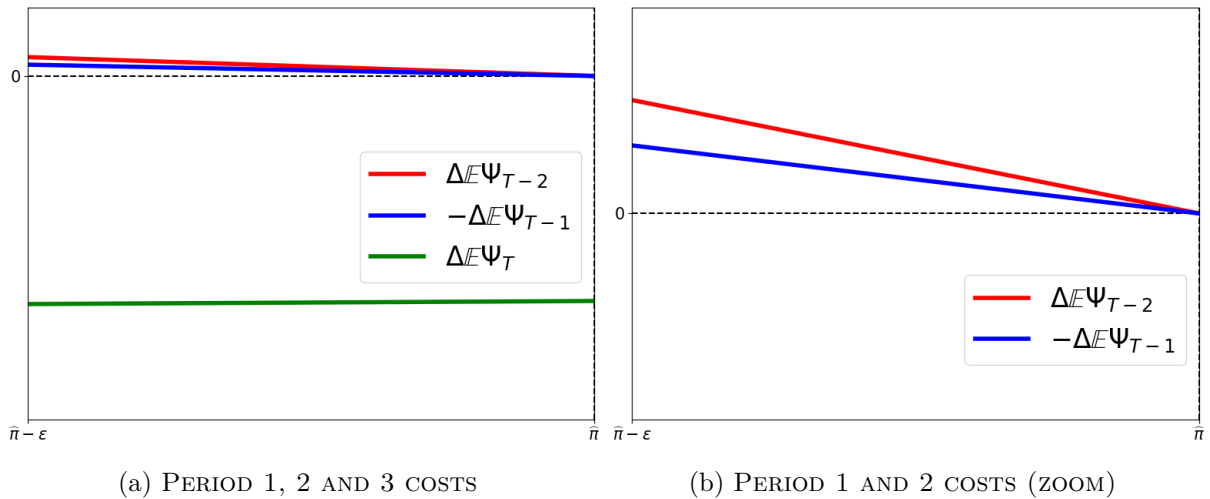
auditor's cost under myopic strategies $(x_{T-2}, y_{T-2} = (0, 1))$ and the cost the auditor deviates to systematic auditing $((x_{T-2}, y_{T-2}) = (1, 1))$.

$$\xi(\pi_{T-2}^0) = \Delta C + \delta \Delta \mathbb{E} \Psi_{T-1} + \delta^2 \Delta \mathbb{E} \Psi_T, \quad \text{where } \pi_{T-2}^0 \rightarrow \hat{\pi}^-.$$

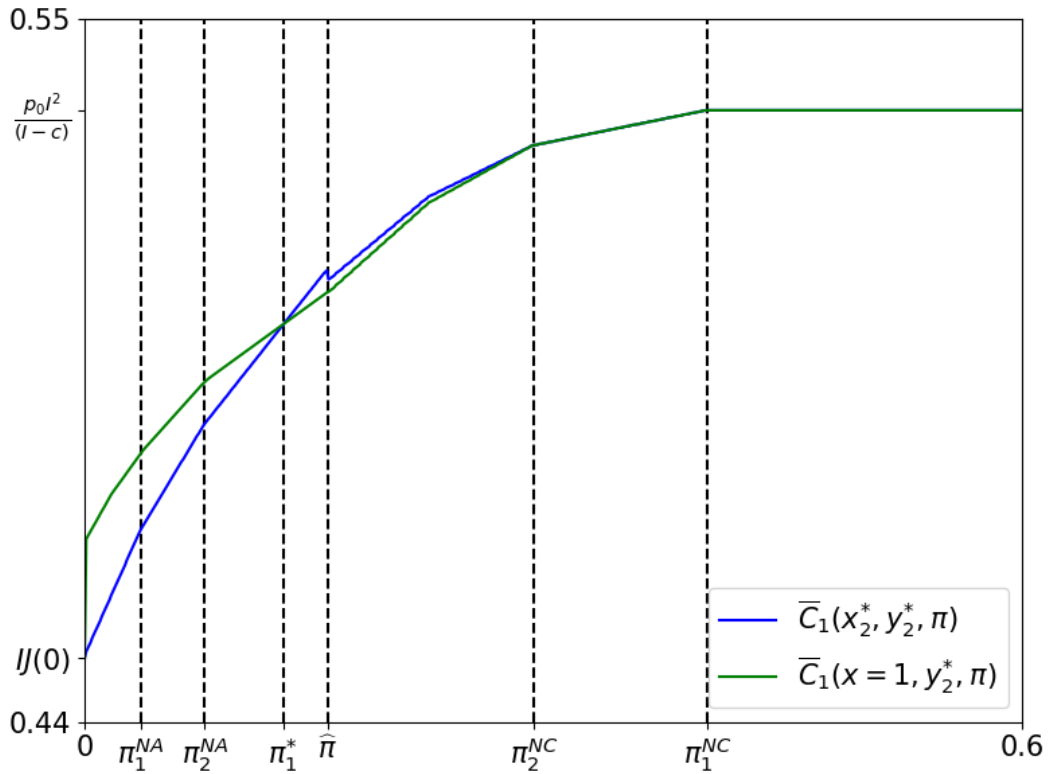
The first term is the difference between both strategies in the instantaneous costs, the second the proceeds from $T - 1$ and the last one the proceeds from T . Figure 4.6 plots each component

to compare them in the vicinity of $\hat{\pi}$. The only term that is positive, i.e., for which a deviation towards systematic auditing is unprofitable, is the current period term. The two other terms are, as shown in the proof, negative. Remembering the definition of $\hat{\pi}$ in equation (4.7), this is the prior that equalizes auditing and not auditing instantaneous costs. Therefore, ΔC tends to 0 when the prior gets closer to $\hat{\pi}$. However this is not sufficient to trigger auditing below $\hat{\pi}$: $\Delta\mathbb{E}\Psi_{T-1}$, the difference in costs between auditing and not auditing in the intermediary period $T-1$ is smaller in absolute value than ΔC (see Figure 4.6b). The $T-1$ term, as a term that account for information gains, is negative, but also tends towards 0 when approaching $\hat{\pi}$. This is why there is no reputation-based deterrence effect below $\hat{\pi}$ at $T=2$. Finally, including a third period adds the $\Delta\mathbb{E}\Psi_T$ term to the balance. This term is negative, but does not tend to 0 at $\hat{\pi}$ (see green line in Figure 4.6a).

Figure 4.6 – MULTI-PERIOD COSTS IN THE VICINITY OF $\hat{\pi}$



Let us again think about the situation under scrutiny through the lens of deviation towards total auditing, i.e., $x = 1$. Figure 4.7 plots the period 1 average discounted cost under the second period equilibrium actions (x_2^*, y_2^*) . It also plots the period 1 average discounted cost under a deviation from the previous action consisting in systematically auditing. We see that this deviation becomes profitable above π_1^* , hence the reputation-based deterrence effect. Note also the discontinuity of the blue curve, i.e., under a myopic equilibrium, at the prior $\hat{\pi}$. This is because auditing, which means resorting with probability $x_1^* > 0$ to transitions φ_{AI} and φ_{AV} instead of φ_{NA} , becomes profitable as soon as $\pi > \pi_1^*$, but does not come into play until $\hat{\pi}$. When it does all of a sudden, the related cost saving makes the cost curve discontinuous.

Figure 4.7 – $\bar{C}_1(x_1, y_1, \pi)$ FOR $T = 3$


4.3.4 Multiperiod setting

In a multiperiod setting where $T \geq 3$, Proposition 4.5 shows that if the horizon is far enough and the insurer is patient enough, reputation-based fraud deterrence in the extensive margin takes place.

Proposition 4.5. *For an arbitrary number of periods T , there is reputation-based deterrence if the remaining relationship is long enough, i.e., for t such that*

$$(1 - p_0 - p_H) \frac{\delta(1 - \delta^{T-t})}{1 - \delta} > 1,$$

which is true for $T - t$ large enough, provided the auditor is patient enough

$$\delta > \frac{1}{2 - p_0 - p_H}.$$

In words, for δ large enough, a type D SP whose reputation is just below the threshold $\hat{\pi}$ will have an incentive to deviate from the instantaneous equilibrium $\hat{y}(\hat{\pi} - \epsilon) = 1$ if $T - t$ is large enough. Intuitively, if the auditor never audits in the neighborhood just below $\hat{\pi}$ and expects a type D to systematically defraud, then not observing a claim leads him to believe once and for

all that the scrutinized auditee is of type H . In this context, by deviating, an auditee of type D will masquerade as a type H for the remaining of the relationship and defraud systematically with no fear of being audited. These future accumulated profits are increasing with the number of remaining periods and the discounting factor, which explains the existence of thresholds in the proposition. In other words, the threat of a future audit deters fraud at the beginning of the relationship.

4.4 Auditing with Deterrence and Learning Effects and the Restless Bandit Problem

In this section, we provide a short description of the relationship between our auditing problem and the restless bandit problem. We leave a formal resolution of the problem for future work. We also set aside the strategic arms aspect.

A fundamental assumption in the regular bandit literature is that the reward associated to each arm only evolves if the arm is chosen. In [Aboutajdine and Picard \(2019b\)](#), this was true because not auditing had an updating function equal to the identity function. In this paper, not auditing corresponds to the updating functions φ_{NA} and φ_{NC} , which means that an arm (an SP) not played (audited) still sees the associated belief evolve. This is due to the difference in claim submission between types. A type D submits valid and mistakenly invalid claims with the same probability as a type H , but the type D also voluntarily submits invalid claims, which means he submits more claims on average than the type H . Therefore, even without auditing, observing a claim indicates that an SP is more likely to be of type D .

This is an important aspect as we can no longer rely on an optimal stopping formulation of the problem. Even if the belief drops below the threshold π_t^* , auditing does not stop and may resume if the belief goes back above the threshold through non auditing based updating.

Figure 4.8 – BELIEF TRAJECTORIES FOR A TYPE D UNDER MYOPIC STRATEGIES

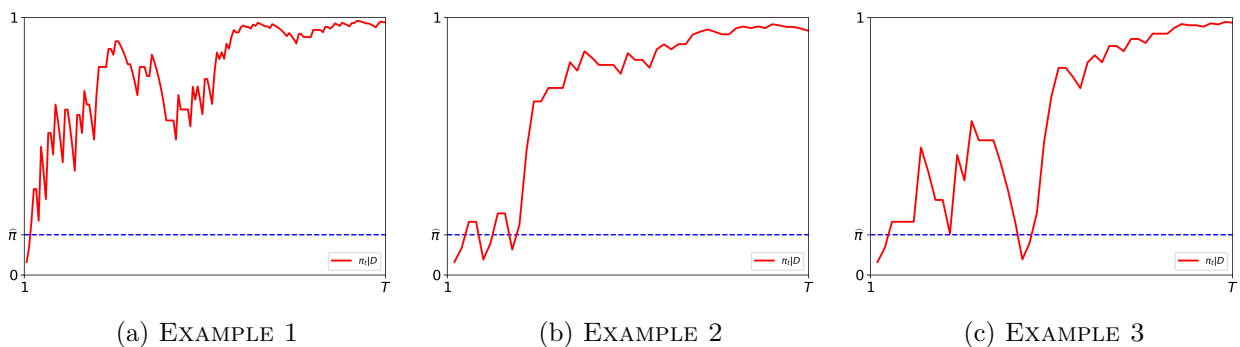
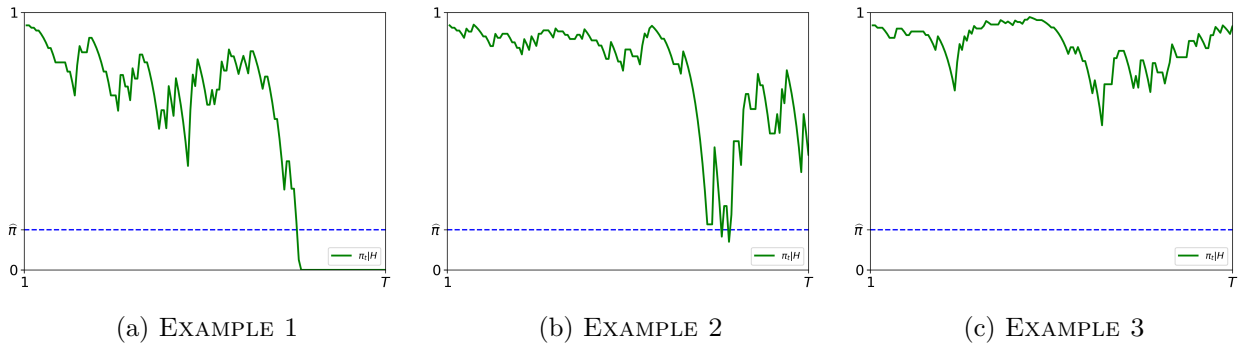


Figure 4.9 – BELIEF TRAJECTORIES FOR A TYPE H UNDER MYOPIC STRATEGIES


Figures 4.8 and 4.9 show simulations of the trajectories of beliefs for a type D and a type H under myopic strategies at all periods, but with updating. Although not realistic, as an insurer who updates beliefs cannot plausibly use myopic strategies, these examples still illustrate the ideas above. In Figure 4.8a, the belief converges to the true one, even after starting below the threshold. In addition, beliefs in Figures 4.8b and 4.8c do cross the threshold more than once before converging to the true belief. In the case of a type H , once he reaches the zone where no auditing is conducted below the threshold, it is sufficient for him to not claim once to see his belief become equal to 0, the true belief. However, as shown by the random draw in Figure 4.9b, if a claim is submitted right after entering the no auditing zone, the SP's belief remains strictly above 0. In addition, these trajectories are obtained under myopic strategies when type D SPs and the insurer do not adapt their behavior earlier in the relationship. Therefore, there is no guarantee that the true beliefs 0 and 1 will be the only stationary points asymptotically.

Therefore, our dynamic auditing problem with deterrence and learning is related to a complex class of bandit problems where the reward distribution of non-played arms still evolves and arms adapt to the player's strategy.

4.5 Conclusion

This paper has analyzed the insurance fraud problem when the auditor faces strategic auditees with a reputation. The objectives of auditing are threefold: recovering illegitimate claims, deterring fraud and learning about auditees' types, and their effects overlap. The main result of our paper is that the combination of information learning by auditing and strategic behavior of auditees produces a reputation-based deterrence effect as a result of their interaction.

An important but ambitious extension of our analysis would be to approach it as a bandit problem, but, as discussed in the last section of the paper, the problem is particularly complex.

One possibility would be to decompose the problem and start with an intermediary setting. Going back to the model of [Aboutajdine and Picard \(2019b\)](#), one could relax only the assumption of having both types submit exactly one claim per period, and keep auditees non-strategic. Then, our problem would be a restless bandit problem, without strategic arms.

4.6 Appendix

4.6.1 Proof of Proposition 4.1

Proof. Under an audit strategy $x_1(\cdot)$, the best response $y_1(\pi_1)$ of a type D SP under insurer's belief π_1 is obtained by maximizing $\Pi(x_1(\pi_1), y_1, \pi_1)$ w.r.t. $y_1 \in [0, 1]$, which gives

$$y_1(\pi_1) = 1 \quad \text{if } (I - \omega)[1 - x_1(\pi_1)] - \omega x_1(\pi_1) > 0, \quad (4.15)$$

$$y_1(\pi_1) \in [0, 1] \quad \text{if } (I - \omega)[1 - x_1(\pi_1)] - \omega x_1(\pi_1) = 0, \quad (4.16)$$

$$y_1(\pi_1) = 0 \quad \text{if } (I - \omega)[1 - x_1(\pi_1)] - \omega x_1(\pi_1) < 0. \quad (4.17)$$

Under a fraud strategy $y_1(\cdot)$, the best response $x_1(\pi_1)$ of the insurer under belief π_1 is obtained by minimizing $C(x_1, y_1(\pi_1), \pi_1)$ w.r.t. $x_1 \in [0, 1]$, which gives

$$x_1(\pi_1) = 1 \quad \text{if } c[p_0 + p_H + \pi_1 y_1(\pi_1)(1 - p_0 - p_H)] - I[p_H + \pi_1 y_1(\pi_1)(1 - p_0 - p_H)] < 0 \quad (4.18)$$

$$x_1(\pi_1) \in [0, 1] \quad \text{if } c[p_0 + p_H + \pi_1 y_1(\pi_1)(1 - p_0 - p_H)] - I[p_H + \pi_1 y_1(\pi_1)(1 - p_0 - p_H)] = 0, \quad (4.19)$$

$$x_1(\pi_1) = 0 \quad \text{if } c[p_0 + p_H + \pi_1 y_1(\pi_1)(1 - p_0 - p_H)] - I[p_H + \pi_1 y_1(\pi_1)(1 - p_0 - p_H)] > 0. \quad (4.20)$$

If $x_1(\pi_1) = 1$, then (4.17) gives $y_1(\pi_1) = 0$. Using (4.3) and (4.20) yields $x_1(\pi_1) = 0$, hence a contradiction. Thus, we have $x_1(\pi_1) < 1$. In other words, systematic auditing cannot be an equilibrium outcome.

If $x_1(\pi_1) \in (0, 1)$, then (4.16) gives $y_1(\pi_1) = \hat{y}(\pi_1)$, with $\hat{y}(\pi_1) \in (0, 1)$ if $\pi_1 > \hat{\pi}$, and $\hat{y}(\pi_1) = 1$ if $\pi_1 \leq \hat{\pi}$. Using (4.15) and (4.16) gives $x_1(\pi_1) = \hat{x} \in (0, 1)$ when $\hat{y}(\pi_1) \in (0, 1)$, i.e., when $\pi_1 > \hat{\pi}$, and $x_1(\pi_1) \leq \hat{x}$ when $\hat{y}(\pi_1) = 1$, i.e., when $\pi_1 = \hat{\pi}$. Hence, there exists a mixed equilibrium where no player audits or defrauds with certainty, such that

$$x_1^*(\pi_1) = \hat{x} \in (0, 1) \text{ and } y_1^*(\pi_1) = \hat{y}(\pi_1) \in (0, 1) \text{ if } \pi_1 > \hat{\pi},$$

Finally, if $x_1(\pi_1) = 0$, then (4.15) gives $y_1(\pi_1) = 1$, and (4.20) is satisfied if $\pi_1 \leq \hat{\pi}$. Hence

there exists an equilibrium such that

$$x_1^*(\pi_1) = 0 \text{ and } y_1^*(\pi_1) = 1 \text{ if } \pi_1 \leq \hat{\pi}.$$

Patching up these cases yields the characterization provided in Proposition 4.1. □

4.6.2 Joint proof of Propositions 4.2 and 4.3

Proof. Assume $\bar{\pi} < \hat{\pi}$. Let $\pi_{AI}(\pi_1, y_1)$, $\pi_{NA}(\pi_1, y_1)$ and $\pi_{NC}(\pi_1, y_1)$ defined by

$$\varphi_{AI}(\pi_{AI}, y_1) = \hat{\pi},$$

$$\varphi_{NA}(\pi_{NA}, y_1) = \hat{\pi},$$

$$\varphi_{NC}(\pi_{NC}, y_1) = \hat{\pi},$$

with

$$\pi_{AI}(\pi_1, y_1) < \pi_{NA}(\pi_1, y_1) < \hat{\pi} < \pi_{NC}(\pi_1, y_1),$$

for all π_1, y_1 .

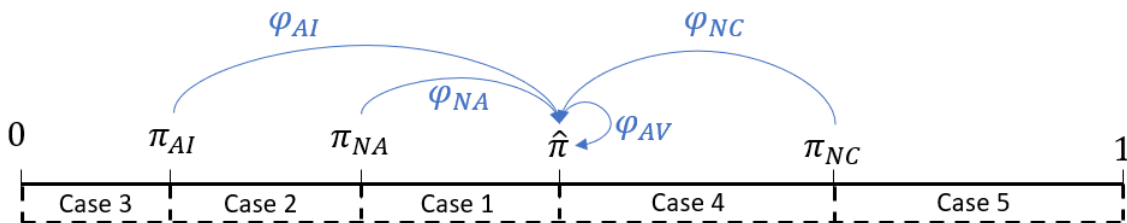
Consider period 1 equilibrium strategies $x_1^*(\pi_1)$ and $y_1^*(\pi_1)$. Let us denote $K^*(\pi_1) \equiv K(\pi_1, y_1^*(\pi_1))$ and

$$\Delta = \frac{p_0 I^2}{(I - c)} - IJ(\bar{\pi}),$$

with $\Delta > 0$ from $\bar{\pi} < \hat{\pi}$.

According to initial beliefs π_1 , five possible cases have to be distinguished.

Figure 4.10 – PROOF CASES



Case 1 : π_1 is such that $\pi_{NA}(\pi_1, y_1^*(\pi_1)) < \pi_1 < \hat{\pi}$.

In the case of an audit when a claim is filed, we have

$$\pi_2 = \pi_1 < \hat{\pi} \quad \text{with prob. } \frac{p_0}{p_0 + K^*(\pi_1)}, \quad (4.21)$$

$$\pi_2 = \bar{\pi} < \hat{\pi} \quad \text{with prob. } \frac{\alpha\pi_1 y_1^*(\pi_1)(1 - p_0 - p_H)}{p_0 + K^*(\pi_1)}, \quad (4.22)$$

$$\pi_2 = \varphi_{AI}(\pi_1, y_1^*(\pi_1)) > \hat{\pi} \quad \text{with prob. } \frac{K^*(\pi_1) - \alpha\pi_1 y_1^*(\pi_1)(1 - p_0 - p_H)}{p_0 + K^*(\pi_1)}. \quad (4.23)$$

If there is no audit, we have

$$\pi_2 = \varphi_{NA}(\pi_1, y_1^*(\pi_1)) > \hat{\pi}. \quad (4.24)$$

The insurer chooses x_1 in $[0, 1]$ in order to minimize

$$\begin{aligned} C(x_1, y_1^*(\pi_1), \pi_1) + \delta \mathbb{E}_{\tilde{\pi}_2} [C(x_2^*(\tilde{\pi}_2), y_2^*(\tilde{\pi}_2), \tilde{\pi}_2)] \\ = C(x_1, y_1^*(\pi_1), \pi_1) + \delta \mathbb{E}_{\tilde{\pi}_2} C(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2). \end{aligned}$$

Using (4.9) and (4.21)-(4.24) shows that x_1 minimizes

$$\begin{aligned} I + x_1 \left[c - I \frac{K^*(\pi_1)}{p_0 + K^*(\pi_1)} \right] + \delta x_1 \left[\frac{p_0}{p_0 + K^*(\pi_1)} IJ(\pi_1) \right. \\ \left. + \frac{\alpha\pi_1 y_1^*(\pi_1)(1 - p_0 - p_H)}{p_0 + K^*(\pi_1)} IJ(\bar{\pi}) + \frac{K^*(\pi_1) - \alpha\pi_1 y_1^*(\pi_1)(1 - p_0 - p_H)}{p_0 + K^*(\pi_1)} \times \frac{p_0 I^2}{I - c} \right] \\ + \delta(1 - x_1) \frac{p_0 I^2}{I - c}, \end{aligned}$$

in $[0, 1]$. Since the function to be minimized is linear w.r.t. x_1 , we deduce

$$\begin{aligned} x_1 \begin{cases} = 1 \\ \in [0, 1] \\ = 0 \end{cases} \quad \text{if } [p_0 + K^*(\pi_1)]c - IK^*(\pi_1) + \delta p_0 \left[J(\pi_1)I - \frac{p_0 I^2}{I - c} \right] \\ \begin{aligned} &< 0 \\ &= 0 \\ &> 0 \end{aligned} \quad . \end{aligned}$$

This gives

$$x_1 \begin{cases} = 1 & > \\ \in [0, 1] & \text{if } y_1^*(\pi_1) = f(\pi_1) \\ = 0 & < \end{cases} \quad (4.25)$$

where

$$f(\pi_1) = \frac{k}{\pi_1} [\hat{\pi} - \eta(\hat{\pi} - \pi_1)], \quad (4.26)$$

and

$$\begin{aligned} k &= \frac{I - c}{I - c + \delta\alpha\Delta} \in (0, 1], \\ \eta &= \frac{\delta p_0 I}{I - c} \in (0, 1) \end{aligned}$$

with $f'(\pi_1) < 0$ and $f(\hat{\pi}) < 1$ if $\alpha > 0$ and $f(\hat{\pi}) = 1$ if $\alpha = 0$, because $k < 1$ when $\alpha > 0$ and $k = 1$ when $\alpha = 0$. Note also that $f(\pi_{**}) = 1$ with

$$\pi_{**} = \hat{\pi} \frac{I - c - \delta p_0 I}{I - c - \delta p_0 I + \delta\alpha\Delta} = \hat{\pi} \frac{k(1 - \eta)}{k(1 - \eta) + 1 - k}, \quad (4.27)$$

and $\pi_{**} < \hat{\pi}$ if $\alpha > 0$, and $\pi_{**} = \hat{\pi}$ if $\alpha = 0$.

Consider now the optimal choice of a type D SP (in the case where he has the opportunity to certify a fraudulent claim, which occurs with probability $1 - p_0 - p_H$). Let $\mathbb{E}_F [\Pi(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2)]$ and $\mathbb{E}_N [\Pi(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2)]$ be the period 2 expected fraud profit of a type D SP in case of fraud or no-fraud at period 1 respectively. Here $\tilde{\pi}_2$ denotes period 2 beliefs in each case (fraud F and no fraud N), and these beliefs depend on initial beliefs π_1 , on the type D SP strategy $y_1^*(\pi_1)$, and on s_1 . For beliefs and equilibrium strategies such that $\pi_{NA}(\pi_1, y_1^*(\pi_1)) < \pi_1 < \hat{\pi}$, if the SP defrauds, we have

$$\tilde{\pi}_2 = \begin{cases} \varphi_{AI}(\pi_1, y_1^*(\pi_1)) > \hat{\pi} & \text{in the case of an audit without} \\ & \text{unambiguous proof of fraud,} \\ \varphi_{NA}(\pi_1, y_1^*(\pi_1)) > \hat{\pi} & \text{if there is no audit.} \end{cases}$$

In both case, the SP's expected profit at period 2 is zero. This is also the case, if an audit provides an unambiguous proof of fraud. Hence, we have

$$\mathbb{E}_F [\Pi(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2)] = 0.$$

If the type D SP does not defraud, no claim is filed and thus, we have

$$\tilde{\pi}_2 = \varphi_{NC}(\pi_1, y_1^*(\pi_1)) < \hat{\pi},$$

which gives

$$\mathbb{E}_N \left[\Pi(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2) \right] = (1 - p_0 - p_H)(I - \omega).$$

Hence, at period 1, the type D best response is given by maximizing

$$y_1 \left\{ [1 - x_1^*(\pi_1)]I - x_1^*(\pi_1)\omega + \delta \mathbb{E}_F \left[\Pi(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2) \right] - \delta \mathbb{E}_N \left[\Pi(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2) \right] \right\},$$

w.r.t. y_1 in $[0, 1]$, which gives

$$y_1 \begin{cases} = 1 & < 0 \\ \in [0, 1] & \text{if } [1 - x_1^*(\pi_1)](I - \omega) - x_1^*(\pi_1)\omega - \delta(1 - p_0 - p_H)(I - \omega) = 0, \\ = 0 & > 0 \end{cases}$$

or equivalently

$$y_1 \begin{cases} = 1 & < \\ \in [0, 1] & \text{if } x_1^*(\pi_1) = x_{**}, \\ = 0 & > \end{cases} \quad (4.28)$$

with

$$x_{**} = \frac{(1 - \delta(1 - p_0 - p_H))(I - \omega)}{I} \in (0, 1). \quad (4.29)$$

Suppose $x_1^*(\pi_1) = 1$. Then (4.25) gives $y_1^*(\pi_1) \geq f(\pi_1) > 0$, and (4.39) implies $x_1^*(\pi_1) \leq x^{**} < 1$, hence a contradiction. If $x_1^*(\pi_1) = 0$, then (4.25) gives $y_1^*(\pi_1) \leq f(\pi_1) < 1$, and (4.39) gives $x_1^*(\pi_1) \geq x^{**} > 0$, which once again is a contradiction. Hence, we have $x_1^*(\pi_1) \in (0, 1)$, and (4.25) implies $y_1^*(\pi_1) = f(\pi_1)$, with $\pi_{**} \leq \pi_1 < \hat{\pi}$. When $\pi_{**} < \pi_1 < \hat{\pi}$, we have $y_1^*(\pi_1) = f(\pi_1) \in (0, 1)$, and (4.39) gives $x_1^*(\pi_1) = x_{**} \in (0, 1)$. When $\pi_1 = \pi_{**}$, we have $y_1^*(\pi_1) = f(\pi_1) = 1$, and (4.39) gives $x_1^*(\pi_1) \leq x_{**}$.

Such an equilibrium exists if $\varphi_{NA}(\pi_1, y_1^*(\pi_1)) < \pi_1 < \hat{\pi}$ or, equivalently, if $\varphi_{NA}(\pi_1, y_1^*(\pi_1)) > \hat{\pi}$ and $\pi_1 < \hat{\pi}$. This is equivalent to

$$\pi_{NA}^* < \pi_1 < \hat{\pi},$$

where

$$\pi_{NA}^* = \hat{\pi} \left[1 - \frac{k(1 - \hat{\pi})(1 - p_0 - p_H)}{p_0 + p_H + k\eta(1 - \hat{\pi})(1 - p_0 - p_H)} \right] < \hat{\pi}.$$

Hence, an equilibrium exists with $x_1^*(\pi_1) = x_{**} \in (0, 1)$ and $y_1^*(\pi_1) = f(\pi_1) \in (0, 1)$ if $\sup\{\pi_{NA}^*, \pi_{**}\} < \pi_1 < \hat{\pi}$, which requires $\alpha > 0$. Everything else given, we have $\pi_{NA}^* < \pi_{**}$, when k is small enough, which corresponds to α small. If $\alpha = 0$ and $x_1^*(\pi_1) > 0$, we would have $y_1^*(\pi_1) \geq f(\pi_1) > 0$ if $\pi_1 < \hat{\pi}$, hence a contradiction. Thus, we have $x_1^*(\pi_1) = 0$ and $y_1^*(\pi_1) = 1$ when $\alpha = 0$.

Case 2 : π_1 is such that $\pi_{AI}(\pi_1, y_1^*(\pi_1)) < \pi_1 < \pi_{NA}(\pi_1, y_1^*(\pi_1))$.

In the case of an audit when a claim is filed, we have

$$\pi_2 = \pi_1 < \hat{\pi} \text{ with prob. } \frac{p_0}{p_0 + K^*(\pi_1)}, \quad (4.30)$$

$$\pi_2 = \bar{\pi} < \hat{\pi} \text{ with prob. } \frac{\alpha\pi_1 y_1^*(\pi_1)(1 - p_0 - p_H)}{p_0 + K^*(\pi_1)}, \quad (4.31)$$

$$\begin{aligned} \pi_2 &= \varphi_{AI}(\pi_1, y_1^*(\pi_1)) > \hat{\pi} \\ &\text{with prob. } \frac{K^*(\pi_1) - \alpha\pi_1 y_1^*(\pi_1)(1 - p_0 - p_H)}{p_0 + K^*(\pi_1)}. \end{aligned} \quad (4.32)$$

If there is no audit, we have

$$\pi_2 = \varphi_{NA}(\pi_1, y_1^*(\pi_1)) < \hat{\pi}. \quad (4.33)$$

In this case, the insurer chooses x_1 in $[0, 1]$ in order to minimize

$$\begin{aligned} &I + x_1 \left[c - I \frac{K^*(\pi_1)}{p_0 + K^*(\pi_1)} \right] + \delta x_1 \left[\frac{p_0}{p_0 + K^*(\pi_1)} IJ(\pi_1) \right. \\ &\left. + \frac{\alpha\pi_1 y_1^*(\pi_1)(1 - p_0 - p_H)}{p_0 + K^*(\pi_1)} IJ(\bar{\pi}) + \frac{K^*(\pi_1) - \alpha\pi_1 y_1^*(\pi_1)(1 - p_0 - p_H)}{p_0 + K^*(\pi_1)} \times \frac{p_0 I^2}{I - c} \right] \\ &\delta(1 - x_1) IJ(\varphi_{NA}(\pi_1, y_1^*(\pi_1))), \end{aligned}$$

in $[0, 1]$. We deduce

$$x_1 \begin{cases} = 1 \\ \in [0, 1] \\ = 0 \end{cases} \quad \text{if } [p_0 + K^*(\pi_1)]c - IK^*(\pi_1) + \delta I[p_H(1 - p_0 - p_H) - K^*(\pi_1)\frac{I(1 - p_0) - c}{I - c} - \alpha\Delta\pi_1 y_1^*(\pi_1)(1 - p_0 - p_H)] \begin{matrix} < 0 \\ = 0 \\ > 0 \end{matrix} .$$

This gives after simplification

$$x_1 \begin{cases} = 1 \\ \in [0, 1] \\ = 0 \end{cases} \quad \text{if } \begin{matrix} > \\ y_1^*(\pi_1) = g(\pi_1) \\ < \end{matrix} \quad (4.34)$$

where

$$g(\pi_1) = \frac{k}{\pi_1} \left[\frac{p_0 \hat{\pi} + \eta p_H (\hat{\pi} - \pi_1)}{p_0 + k\eta(1 - p_0 - p_H)(1 - \hat{\pi})} \right], \quad (4.35)$$

with $g'(\pi_1) < 0$. Note that

$$f(\pi_{NA}^*) = g(\pi_{NA}^*) = \frac{k\hat{\pi}(p_0 + p_H)}{p_0 + p_H - k(1 - \eta)(1 - \hat{\pi})(1 - p_0 - p_H)},$$

which corresponds to the continuity between cases 1 and 2 at $\pi_1 = \pi_{NA}^*$. Reasoning as in Case 1 shows that when $\alpha > 0$ there exists an equilibrium in mixed strategies $x_1^*(\pi_1) = x_{***} \in (0, 1)$ and $y_1^*(\pi_1) = g(\pi_1) \in (0, 1)$ if $\sup\{\pi_{AI}^*, \pi_{***}\} < \pi_1 < \pi_{NA}^*$ and $\pi_{NA}^* < \pi_{**}$, where π_{AI}^* and x_{***} are defined by $\pi_{AI}(\pi_1, y_1^*(\pi_1)) = \hat{\pi}$, or equivalently $\varphi_{AI}(\pi_1, y_1^*(\pi_1)) > \hat{\pi}$, and by $g(\pi_{***}) = 1$, respectively. x_{***} is the audit probability that makes the type D SP indifferent between defrauding or not-defrauding.

Case 3 : π_1 is such that $\pi_1 < \pi_{AI}(\pi_1, y_1)$.

Straight forward calculation shows that $x_1^*(\pi_1) = 0$ and $y_1^*(\pi_1) = 1$ in that case. This case holds when $\pi_1 < \sup\{\pi_{AI}^*, \pi_{***}\}$.

Case 4 : $\hat{\pi}$ is such that $\hat{\pi} < \pi_1 < \pi_{NC}(\pi_1, y_1^*(\pi_1))$.

When $\alpha = 0$, in the case of an audit when a claim is filed, we have

$$\pi_2 = \pi_1 > \hat{\pi} \quad \text{with prob. } \frac{p_0}{p_0 + K^*(\pi_1)}, \quad (4.36)$$

$$\pi_2 = \varphi_{AI}(\pi_1, y_1^*(\pi_1)) > \hat{\pi} \quad \text{with prob. } \frac{K^*(\pi_1)}{p_0 + K^*(\pi_1)}. \quad (4.37)$$

If there is no audit, we have

$$\pi_2 = \varphi_{NA}(\pi_1, y_1^*(\pi_1)) > \hat{\pi}. \quad (4.38)$$

In this case the type D chooses y_1 in $[0, 1]$ in order to maximize

$$y_1 \left\{ [1 - x_1^*(\pi_1)]I - x_1^*(\pi_1)\omega + \delta \mathbb{E}_F [\Pi(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2)] - \delta \mathbb{E}_N [\Pi(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2)] \right\},$$

w.r.t. y_1 in $[0, 1]$, which gives

$$y_1 \begin{cases} = 1 & < 0 \\ \in [0, 1] & \text{if } [1 - x_1^*(\pi_1)](I - \omega) - x_1^*(\pi_1)\omega - \delta(1 - p_0 - p_H)(I - \omega) = 0, \\ = 0 & > 0 \end{cases}$$

or equivalently

$$y_1 \begin{cases} = 1 & < \\ \in [0, 1] & \text{if } x_1^*(\pi_1) = \hat{x}, \\ = 0 & > \end{cases} \quad (4.39)$$

with

$$\hat{x} = \frac{(I - \omega)(1 - \delta(1 - p_0 - p_H))}{I} = (1 - \delta(1 - p_0 - p_H))\hat{\pi} \in (0, 1). \quad (4.40)$$

Consider now the optimal choice of the insurer. She has to minimize

$$C(x_1, y_1^*(\pi_1), \pi_1) + \delta \mathbb{E}_{\tilde{\pi}_2} [C(x_2^*(\tilde{\pi}_2), y_2^*(\tilde{\pi}_2), \tilde{\pi}_2)] = C(x_1, y_1^*(\pi_1), \pi_1) + \delta \left[\frac{p_0 I^2}{I - c} \right].$$

The second period cost is independent of x_1 . Indeed, all outcomes NA, AV and AI lead to posteriors larger than $\hat{\pi}$. Therefore, the insurer is made indifferent for the myopic fraud level $\hat{y}(\pi_1) = \frac{\hat{\pi}}{\pi_1} \in (0, 1)$.

For $\alpha = 0$, the equilibrium in Case 4 is then given by

$$x_1^*(\pi_1) = \frac{(I - \omega)(1 - \delta(1 - p_0 - p_H))}{I} < \hat{x} \quad \text{and} \quad y_1(\pi_1) = \frac{\hat{\pi}}{\pi_1}.$$

Case 5 : $\hat{\pi}$ is such that $\pi_1 > \pi_{NC}(\pi_1, y_1^*(\pi_1))$.

In this case, all priors in period 2 obtained from π_1 through updating are larger than $\hat{\pi}$. Then the expectation of period 2 cost is known and independent of x_1

$$\mathbb{E}_{\tilde{\pi}_2} [C(x_2^*(\tilde{\pi}_2), y_2^*(\tilde{\pi}_2), \tilde{\pi}_2)] = \frac{p_0 I^2}{I - c}.$$

The type D has to maximize

$$y_1 \left\{ [1 - x_1^*(\pi_1)]I - x_1^*(\pi_1)\omega + \delta \mathbb{E}_F \left[\Pi(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2) \right] - \delta \mathbb{E}_N \left[\Pi(\hat{x}(\tilde{\pi}_2), \hat{y}(\tilde{\pi}_2), \tilde{\pi}_2) \right] \right\},$$

where the period 2 terms are all equal to zero since all updated beliefs are larger than $\hat{\pi}$. Therefore, the period 1 equilibrium audit and fraud are the myopic ones, for beliefs larger than $\pi_{NC}(\pi_1, y_1^*(\pi_1))$

$$x_1^*(\pi_1) = \hat{x} = \frac{I - \omega}{I} \quad \text{and} \quad y_1^*(\pi_1) = \hat{y}(\pi_1) = \frac{\hat{\pi}}{\pi_1}.$$

□

4.6.3 Proof of Proposition 4.4

Proof. Let $T = 3$. We already know that the proposition is valid when $\alpha > 0$, and thus we may restrict the proof to the case $\alpha = 0$. Let

$$\Psi(\pi) = \inf\{IJ(\pi), p_0I^2/(I - c)\}.$$

$\Psi(\pi)$ is the expected cost of the insurer at the equilibrium of the instantaneous game, under beliefs π . Hence,

$$C(x_T^*(\pi_T), y_T^*(\pi_T), \pi_T) = \Psi(\pi_T),$$

for all π_T in $[0, 1]$, since T is the last period of the game. Furthermore, when $\alpha = 0$, the equilibrium strategies of period $T - 1$ also coincide with those of the instantaneous game, with

$$C(x_{T-1}^*(\pi_{T-1}), y_{T-1}^*(\pi_{T-1}), \pi_{T-1}) = \Psi(\pi_{T-1}),$$

for all π_{T-1} in $[0, 1]$.

Assume

$$x_{T-2}^*(\pi_{T-2}) = \hat{x}(\pi_{T-2}), \tag{4.41}$$

$$y_{T-2}^*(\pi_{T-2}) = \hat{y}(\pi_{T-2}), \tag{4.42}$$

for all $\pi_{T-2} \in [0, 1]$. Let $\pi_{T-2}^0 \in (0, \hat{\pi})$, π_{T-2}^0 close to $\hat{\pi}$. (4.41) and (4.42) imply

$$\begin{aligned} x_{T-2}^*(\pi_{T-2}^0) &= 0, \\ y_{T-2}^*(\pi_{T-2}^0) &= 1. \end{aligned}$$

Consider a deviation of the insurer's strategy from $x_{T-2}^*(\pi_{T-2}^0) = 0$ to $x_{T-2}(\pi_{T-2}^0) = 1$. Let Ψ_{T-1}^N and Ψ_{T-1}^A denote the insurer's expected cost at period $T-1$ for beliefs π_{T-2}^0 , if she does not deviate (**N**o audit) and if she deviates (**A**udit) at period $T-2$.

Consider first the case where the insurer does not deviate from her equilibrium strategy $x_{T-2}^*(\pi_{T-2}^0) = 0$. We have

$$\pi_{T-1} = 0 \text{ with prob. } (1 - \pi_{T-2}^0)(1 - p_0 - p_H), \quad (4.43)$$

$$\pi_{T-1} = \varphi_{NA}(\pi_{T-2}^0, 1) > \hat{\pi} \text{ with prob. } 1 - (1 - \pi_{T-2}^0)(1 - p_0 - p_H), \quad (4.44)$$

and thus

$$\begin{aligned} \mathbb{E}\Psi_{T-1}^N &= I(p_0 + p_H)(1 - \pi_{T-2}^0)(1 - p_0 - p_H) \\ &\quad + \frac{p_0 I^2}{I - c} [1 - (1 - \pi_{T-2}^0)(1 - p_0 - p_H)], \end{aligned}$$

where the expected value of the insurer's cost is conditional on period $T-2$ belief π_{T-2}^0 .

When the insurer deviates from her equilibrium strategy to $x_{T-2}(\pi_{T-2}^0) = 1$, we have

$$\pi_{T-1} = 0 \text{ with prob. } (1 - \pi_{T-2}^0)(1 - p_0 - p_H), \quad (4.45)$$

$$\pi_{T-1} = \pi_{T-2}^0 \text{ with prob. } p_0, \quad (4.46)$$

$$\pi_{T-1} = \varphi_{AI}(\pi_{T-2}^0, 1) > \hat{\pi} \text{ with prob. } 1 - (1 - \pi_{T-2}^0)(1 - p_0 - p_H) - p_0, \quad (4.47)$$

and thus

$$\begin{aligned} \mathbb{E}\Psi_{T-1}^A &= I(p_0 + p_H)(1 - \pi_{T-2}^0)(1 - p_0 - p_H) \\ &\quad + I p_0 [p_0 + p_H + \pi_{T-2}^0](1 - p_0 - p_H) \\ &\quad + \frac{p_0 I^2}{I - c} [1 - (1 - \pi_{T-2}^0)(1 - p_0 - p_H) - p_0]. \end{aligned}$$

We deduce

$$\Delta \mathbb{E}\Psi_{T-1} \equiv \mathbb{E}\Psi_{T-1}^A - \mathbb{E}\Psi_{T-1}^N = -p_0 I (1 - p_0 - p_H) (\hat{\pi} - \pi_{T-2}^0). \quad (4.48)$$

Consider now the strategies played at period T following either an audit or no audit at period $T - 2$, and denote Ψ_T^A and Ψ_T^N the corresponding expected costs.

Suppose first that no audit has been performed at period $T - 2$. Consider the two cases that may occur.

If $\pi_{T-1} = 0$, we have $\pi_T = 0$, and thus

$$\mathbb{E}[\Psi_T^N(\pi_T) \mid \pi_{T-1} = 0] = I(p_0 + p_H). \quad (4.49)$$

If $\pi_{T-1} = \varphi_{NA}(\pi_{T-2}^0, 1)$, we have

$$\pi_T \geq \varphi_{s_{T-1}}(\varphi_{NA}(\pi_{T-2}^0, 1), \hat{y}(\varphi_{NA}(\pi_{T-2}^0, 1))), \quad (4.50)$$

in all scenarios s_{T-1} that may occur at period $T - 1$. Note that

$$\varphi_{NA}(\hat{\pi}, 1) = \frac{\hat{\pi}}{p_0 + p_H + \hat{\pi}(1 - p_0 - p_H)}$$

and thus

$$\begin{aligned} \varphi_{s_{T-1}}(\varphi_{NA}(\hat{\pi}, 1), \hat{y}(\varphi_{NA}(\hat{\pi}, 1))) &= \varphi_{s_{T-1}}\left(\varphi_{NA}(\hat{\pi}, 1), \frac{\hat{\pi}}{\varphi_{NA}(\hat{\pi}, 1)}\right) \\ &= \hat{\pi} \frac{1 - p_0 - p_H}{\hat{\pi} + (1 - \hat{\pi})(p_0 + p_H)} \\ &> \hat{\pi}, \end{aligned}$$

where the inequality results from (4.13). Consequently, we have

$$\varphi_{s_{T-1}}(\varphi_{NA}(\pi_{T-2}^0, 1), \hat{y}(\varphi_{NA}(\pi_{T-2}^0, 1))) > \hat{\pi}, \quad (4.51)$$

when π_{T-2}^0 is close to $\hat{\pi}$. Hence, (4.44) and (4.45) show that $\pi_T > \hat{\pi}$ when π_{T-2}^0 is close to $\hat{\pi}$, in all possible scenarios following $\pi_{T-1} = \varphi_{NA}(\pi_{T-2}^0, 1)$.

Consequently, we have

$$\mathbb{E}[\Psi_T^N(\pi_T) \mid \pi_{T-1} = \varphi_{NA}(\pi_{T-2}^0, 1)] = \frac{p_0 I^2}{I - c}, \quad (4.52)$$

and using (4.43),(4.44),(4.49) and (4.52) gives

$$\begin{aligned}\mathbb{E}\Psi_T^N &= (1 - \pi_{T-2}^0)(1 - p_0 - p_H)(p_0 + p_H)I \\ &\quad + [1 - (1 - \pi_{T-2}^0)(1 - p_0 - p_H)] \frac{p_0 I^2}{I - c},\end{aligned}$$

where the expected value of the insurer's cost at period T is conditional on beliefs π_{T-2}^0 .

Suppose now that an audit has been performed at period $T - 2$. Consider the three cases that may occur.

If $\pi_{T-1} = 0$, as in the previous case we have $\pi_T = 0$, and

$$\mathbb{E}[\Psi_T^N(\pi_T) \mid \pi_{T-1} = 0] = I(p_0 + p_H). \quad (4.53)$$

If $\pi_{T-1} = \pi_{T-2}^0$, then using $x_{T-1}^*(\pi_{T-2}^0) = \hat{x}(\pi_{T-2}^0) = 0$ and $y_{T-1}^*(\pi_{T-2}^0) = \hat{y}(\pi_{T-2}^0) = 1$ gives

$$\begin{aligned}\pi_T &= 0 \text{ with prob. } (1 - \pi_{T-2}^0)(1 - p_0 - p_H), \\ \pi_T &= \varphi_{NA}(\pi_{T-2}^0, 1) > \hat{\pi} \text{ with prob. } 1 - (1 - \pi_{T-2}^0)(1 - p_0 - p_H).\end{aligned}$$

If $\pi_{T-1} = \varphi_{AI}(\pi_{T-2}^0, 1)$, we have Δ

$$\pi_T \geq \varphi_{s_i}(\varphi_{AI}(\pi_{T-2}^0, 1), \hat{y}(\varphi_{AI}(\pi_{T-2}^0, 1))), \quad (4.54)$$

in all scenarios s_{T-1} that may occur at period $T - 1$. Since $\varphi_{AI}(\pi_{T-2}^0, 1) > \varphi_{NA}(\pi_{T-2}^0, 1)$, using (4.13) yields $\pi_T > \hat{\pi}$ in all scenarios that may occur at period $T - 1$, following $\pi_{T-1} = \varphi_{AI}(\pi_{T-2}^0, 1)$.

Hence, we have

$$\begin{aligned}\mathbb{E}\Psi_T^A &= (1 - \pi_{T-2}^0)(1 - p_0 - p_H)(p_0 + p_H)I \\ &\quad + p_0 \left[(1 - \pi_{T-2}^0)(1 - p_0 - p_H)(p_0 + p_H)I \right. \\ &\quad \quad \left. + [1 - (1 - \pi_{T-2}^0)(1 - p_0 - p_H)] \frac{p_0 I^2}{I - c} \right] \\ &\quad + [1 - (1 - \pi_{T-2}^0)(1 - p_0 - p_H)] \frac{p_0 I^2}{I - c}.\end{aligned}$$

We obtain

$$\begin{aligned}
 \Delta \mathbb{E}\Psi_T &\equiv \mathbb{E}\Psi_T^A - \mathbb{E}\Psi_T^N = p_0(1 - \pi_{T-2}^0)(1 - p_0 - p_H) \\
 &\quad \times \left[(p_0 + p_H)I - \frac{p_0 I^2}{I - c} \right] \\
 &= -p_0 I \hat{\pi} (1 - p_0 - p_H) (1 - \pi_{T-2}^0). \tag{4.55}
 \end{aligned}$$

Finally, a deviation from $x_{T-2}^*(\pi_{T-2}^0) = 0$ to $x_{T-2}(\pi_{T-2}^0) = 1$ induces a change ΔC in insurance cost at period $T - 2$ defined by

$$\begin{aligned}
 \Delta C &= C(1, 1, \pi_{T-2}^0) - C(0, 1, \pi_{T-2}^0) \\
 &= cp_0 - (I - c)[p_h + (1 - p_0 - p_H)\pi_{T-2}^0],
 \end{aligned}$$

where cp_0 is the expected audit cost for valid claims and $(I - c)[p_h + (1 - p_0 - p_H)\pi_{T-2}^0]$ is the expected value of recouped indemnities, net of audit cost, for detected invalid claims.

Let

$$\xi(\pi_{T-2}^0) = \Delta C + \delta \Delta \mathbb{E}\Psi_{T-1} + \delta^2 \Delta \mathbb{E}\Psi_T,$$

the discounted expected variation in the insurer's cost after the deviation. We have $\Delta C \rightarrow 0$ and $\Delta \mathbb{E}\Psi_{T-1} \rightarrow 0$ when $\pi_{T-2}^0 \rightarrow \hat{\pi}$, and thus $\xi(\pi_{T-2}^0) \rightarrow -\delta^2 p_0 I \hat{\pi} (1 - p_0 - p_H) (1 - \hat{\pi}) < 0$ when $\pi_{T-2}^0 \rightarrow \hat{\pi}$, which shows that the deviation is profitable to the insurer when π_{T-2}^0 is close to $\hat{\pi}$. \square

4.6.4 Proof of Proposition 4.5

Assume there exists $t \geq 1$ such that $y_{t,T}^*(\pi_t) = 1$ for $\pi_t \in [\hat{\pi} - \epsilon, \hat{\pi}] \equiv \mathcal{A}$ where $\epsilon > 0$. Assume an SP of type D with reputation π_t has no legitimate claim to channel and either defrauds with certainty ($y_{t,T} = 1$) or does nothing ($y_{t,T} = 0$). If the SP deviates and $y_{t,T} = 0$, then, by Bayes law, $\pi_{t+1} = 0$ and he will be able to defraud in all subsequent periods without being audited (since $x_{t,T}^* = 0$ at any time t). Denoting $\mathbb{E}_F[V(\tilde{\pi}_{t+1})|\pi_t]$ and $\mathbb{E}_N[V(\tilde{\pi}_{t+1})|\pi_t]$ the time t discounted expected gain over periods $t, t + 1, \dots, T$ of a type D with reputation π_t who respectively defrauds (F, $y_{t,T} = 1$) or does not defraud (N, $y_{t,T} = 0$), for $\pi_t \in [\hat{\pi} - \epsilon, \hat{\pi}]$, we then have

$$\begin{aligned}\mathbb{E}_N[V(\tilde{\pi}_{t+1})|\pi_t] &= (I - \omega)(1 - p_0 - p_H) \sum_{i=1}^{T-t} \delta^i \\ &= (I - \omega)(1 - p_0 - p_H) \frac{\delta(1 - \delta^{T-t})}{1 - \delta},\end{aligned}\tag{4.56}$$

If the type D SP chooses to defraud with $y_{t,T} = 1$, then $\pi_{t+1} = \varphi_{NA}(\pi_t, 1)$ with probability $(1 - x_{t,T}^*(\pi_{t,T}))$ and $\pi_{t+1} = \varphi_{AI}(\pi_t, 1)$ with probability $(x_{t,T}^*(\pi_{t,T}))$. Therefore $\pi_{t+1} > \hat{\pi}$ in both cases for ϵ small enough. Then, at period $t + 1$, the equilibrium is a mixed one and the SP of type D is indifferent between defrauding or not. So the optimal expected gain at time $t + 1$ is obtained by looking at total fraud $y_{t,T} = 1$. In particular, if he defrauds with certainty when given the chance at all subsequent periods, his priors π_{t+i} will always increase and stay above $\hat{\pi}$. Therefore, the SP no longer has any expected gain starting from period $t + 1$ if he defrauds with certainty at period t . Then

$$\begin{aligned}\mathbb{E}_N[V(\tilde{\pi}_{t+1})|\pi_{t,T}] &= (I - \omega)(1 - x_{t,T}^*(\pi_{t,T})) - \omega x_{t,T}^*(\pi_t) \\ &\leq (I - \omega),\end{aligned}\tag{4.57}$$

Equations (4.56) and (4.57) imply that a sufficient condition for

$$\mathbb{E}_N[V(\tilde{\pi}_{t+1})|\pi_t] > \mathbb{E}_F[V(\tilde{\pi}_{t+1})|\pi_t],$$

is given by

$$(1 - p_0 - p_H) \frac{\delta(1 - \delta^{T-t})}{1 - \delta} > 1,$$

which is true for $T - t$ large enough if

$$\delta > \frac{1}{2 - p_0 - p_H}.$$

Conclusion

This thesis tackles the insurance fraud problem through the scope of optimal auditing, specifically when there exist intermediaries between the insurer and its policyholders, and they interact with the insurer on frequent basis.

Chapter 1 focuses on the magnitude of auditing deterrence effects on opticians who might defraud, in the context of the introduction of a fraud detection system. By allowing for cost-effective and better targeted audits, this system turns auditing into a more credible threat. As a consequence, opticians with the most suspicious claims are now audited while it might not have been the case pre-detection system, and their defrauding efforts decrease as a consequence of auditing. This decrease is the strongest for audits who request specific supporting documents, corroborating the idea that commitment and credibility are fundamental ingredients of counter-fraud measures.

Chapter 2 considers auditing as a learning mechanism. When auditor/intermediaries interactions are repeated, the auditor can use information gathered early in the relationship to perform better targeted audits later on. In this framework, optimal auditing strategies take the shape of a threshold in the reputation of intermediaries where reputations worse than the threshold are audited. Earlier in the relationship, this threshold is lower, meaning that more intermediaries are targeted. Furthermore, this threshold is decreasing in the claim value, as, all else being equal, expected proceeds from a successful audit are larger.

Chapter 3 extends Chapter 2's framework to an arbitrarily long relationship. Thresholds are decreasing in the number of remaining periods, and tend to a lower bound threshold when the number of periods goes to infinity. In other words, the further the time horizon, the stronger the learning generated profits from auditing, and the more auditees are monitored at the first stages of the relationship. Auditing is then a mechanism to separate the wheat from the chaff, by progressively excluding honest intermediaries from the auditing pool and isolating dishonest ones for future audits. The auditor's propensity to learn is also increased by polarized behavior

between honest and dishonest auditees, because auditing is more informative.

Finally, Chapter 4 builds on Chapters 2 and 3 and combines both deterrence and learning mechanisms in the same dynamic game theoretic model. Contrary to the previous chapters where learning allowed the auditor to directly generate more auditing proceeds in the future, learning now affects the auditor's optimal strategies indirectly by strengthening the deterrence mechanism, revealing a reputation-based deterrence mechanism. In addition to the threat of being caught for a current defrauding attempt, a dishonest auditee must face the risk of worsening his reputation for future periods. Such a reputation is generated through the learning mechanism of auditing. As a consequence, the auditor audits more at an early period and the dishonest auditee defrauds less. This change in auditing/fraud takes two forms. First, auditing increases in the extensive margin in that auditees with good enough reputations to not be audited in the stage game are now investigated with a non null probability at the early stage. The latter auditees, who would have defrauded with certainty in the stage game, only defraud with probability strictly smaller than one at the early stage. Second, auditing decreases in the intensive margin for auditees with bad enough reputations to be audited in the stage game. Because of this reputation-based threat, the auditor needs to make less auditing efforts to deter dishonest auditees while the latter need to defraud less to discourage the former from systematically auditing.

Future research In Chapter 1, empirical results emphasize the importance of commitment in auditing. Chapter 4 does not take this aspect into account as auditing and fraud decisions are simultaneous. An interesting and straightforward extension would be to consider a Stackelberg game where the auditor has first-mover advantage and can commit to his auditing strategy.

In Chapters 2 to 4, information acquisition played a central role, either as the only mechanism at stake (Chapters 2 and 3), or, more importantly, as a deterrence strengthening mechanism (Chapter 4). This aspect connects this work to the recent and growing literature on Information Design⁹ and suggests several directions to extend the obtained results. For example, introducing information theoretic tools¹⁰ to accurately characterize uncertainty reduction through auditing and how it impacts equilibrium strategies could reveal new insights about the trade-off between exploration and exploitation. In addition, an interesting question that arises in Chapter 3 was to understand how close to the truth did the auditor get. Because, in said chapter, the auditee was non strategic, the auditor ended up learning the truth, or at least learning enough to obtain

⁹See Hörner and Skrzypacz (2017) for a recent review.

¹⁰See Neyman and Okada (1999), Lehrer and Smorodinsky (2000), Zhong (2017).

his optimal payoff. This is not guaranteed in the model from Chapter 4 as the strategic auditee seeks to maintain his reputation good enough. The concept of merging¹¹ captures the idea that a subjective belief becomes closer to the truth when sufficient observations are made, and could help understand how beliefs would behave asymptotically in an infinite horizon version of Chapter 4's model.¹² Another direction would be to reformulate the dynamic game theoretic model from Chapter 4 as a restless bandit problem with strategic arms.¹³ As explained Chapter 4, this class of problems is particularly complex, and an initial, but intermediary step, might be to start by reconsidering the model from Chapter 3 without strategic auditees, but with alternative information acquisition mechanisms other than auditing. Finally, another class of bandit problems that would be of interest is the class of bandits with correlated arms. In this thesis, arms, i.e., auditees, are independent from one another, which excludes the possibility of informational externalities. Yet, if we consider arms to be categories of claims rather than auditees, a natural intuition would be to expect that claims of the same category would be correlated and that, with decreasing informational returns to auditing, it would be optimal to audit a strict subset of claims of the same category, and diversify auditing to include some claims of each category.

Finally, this thesis leaves aside the legal and contractual dimension of the insurance fraud problem. First, learning effects and information acquisition are constrained by privacy laws, for example to guarantee a right to be forgotten.¹⁴ In the case of France, CNIL¹⁵ has issued a Single Authorization¹⁶ governing personal data processing and storage in the context of counter-fraud efforts. One requirement is for suspicious cases who have not been unambiguously proved to be fraudulent to be deleted within six months. Even cases proved to be fraudulent cannot be stored more than five years. Second, fraud can also be tackled through contract design. The insurance academic literature provides several examples of optimal contracts under fraud.¹⁷ In this thesis' setting, one could wonder to which extent can dynamic contract design help mitigate the insurance fraud problem.

¹¹See Blackwell and Dubins (1962), Lehrer and Smorodinsky (1996), Lehrer and Smorodinsky (2000).

¹²See for example Gossner and Tomala (2008) for a bound for the distance between the belief and the truth under costly learning in sequential settings.

¹³Not to be confused with strategic bandits where several players facing the same bandit might collaborate/compete, but the arms are still non-strategic. See Bolton and Harris (1999), Keller, Rady and Cripps (2005), Rosenberg, Solan and Vieille (2007).

¹⁴See GDPR in the EU <https://gdpr-info.eu/issues/right-to-be-forgotten/>

¹⁵Commission nationale de l'informatique et des libertés, <https://www.cnil.fr/en/home>

¹⁶https://www.cnil.fr/sites/default/files/typo/document/FICHE_05_assurance.pdf

¹⁷See for example Bond and Crocker (1997), Crocker and Morgan (1998), Picard (2000).

Bibliography

- Aboutajdine, Reda, and Pierre Picard.** 2018. “Preliminary Investigations for Better Monitoring: Learning in Repeated Insurance Audits.” *Risks*, 6(1): 15.
- Aboutajdine, Reda, and Pierre Picard.** 2019a. “Reputation Effects in Repeated Audits, with Application to Insurance Fraud Deterrence.” Working Paper.
- Aboutajdine, Reda, and Pierre Picard.** 2019b. “Should I Stalk or Should I Go? An Auditing Exploration/Exploitation Dilemma.” Working Paper.
- Advani, Arun, William Elming, and Jonathan Shaw.** 2019. “The Dynamic Effects of Tax Audits.” *Competitive Advantage in the Global Economy (CAGE)* 414.
- Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien.** 1991. “Optimal Learning by Experimentation.” *The Review of Economic Studies*, 58(4): 621–654.
- Allingham, Michael G., and Agnar Sandmo.** 1972. “Income Tax Evasion: A Theoretical Analysis.” *Journal of Public Economics*, 1(3-4): 323–338.
- Allison, Paul D., and Richard P. Waterman.** 2002. “Fixed-Effects Negative Binomial Regression Models.” *Sociological Methodology*, 32(1): 247–265.
- Arrow, K. J., D. Blackwell, and M. A. Girshick.** 1949. “Bayes and Minimax Solutions of Sequential Decision Problems.” *Econometrica*, 17(3/4): 213–244.
- Artís, Manuel, Mercedes Ayuso, and Montserrat Guillén.** 2002. “Detection of Automobile Insurance Fraud With Discrete Choice Models and Misclassified Claims.” *Journal of Risk and Insurance*, 69(3): 325–340.
- Asterès.** 2015. “La filière de la santé visuelle : analyse et proposition.” <https://asteres.fr/etude/filiere-sante-visuelle/>.

- Avenhaus, Rudolf, Bernhard Von Stengel, and Shmuel Zamir.** 2002. “Chapter 51 Inspection Games.” In *Handbook of Game Theory with Economic Applications*. Vol. 3, 1947–1987. Elsevier.
- Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou.** 2006. “Who’s Who in Networks. Wanted: The Key Player.” *Econometrica*, 74(5): 1403–1417.
- Becker, Gary S.** 1968. “Crime and Punishment: An Economic Approach.” *Journal of Political Economy*, 76(2): 169–217.
- Bergemann, Dirk, and Juuso Välimäki.** 2008. “Bandit Problems.” In *The New Palgrave Dictionary of Economics: Volume 1 – 8.*, ed. Steven N. Durlauf and Lawrence E. Blume, 336–340. London:Palgrave Macmillan UK.
- Blackwell, David, and Lester Dubins.** 1962. “Merging of Opinions with Increasing Information.” *The Annals of Mathematical Statistics*, 33(3): 882–886.
- Bolton, Patrick, and Christopher Harris.** 1999. “Strategic Experimentation.” *Econometrica*, 67(2): 349–374.
- Bond, Eric W, and Keith J Crocker.** 1997. “Hardball and the Soft Touch: The Economics of Optimal Insurance Contracts with Costly State Verification and Endogenous Monitoring Costs.” *Journal of Public Economics*, 63(2): 239–264.
- Boning, William C., John Guyton, II Ronald H. Hodge, Joel Slemrod, and Ugo Troiano.** 2018. “Heard It Through the Grapevine: Direct and Network Effects of a Tax Enforcement Field Experiment.” National Bureau of Economic Research Working Paper 24305.
- Borch, Karl.** 1990. *Economics of Insurance*. Elsevier.
- Bourgeon, Jean-Marc, Pierre Picard, and Jerome Pouyet.** 2008. “Providers’ Affiliation, Insurance and Collusion.” *Journal of Banking & Finance*, 32(1): 170–186.
- Boyer, Martin.** 1999. “When Is The Proportion of Criminal Elements Irrelevant? A Study of Insurance Fraud When Insurers Cannot Commit.” In *Automobile Insurance: Road Safety, New Drivers, Risks, Insurance Fraud and Regulation. Hiebner International Series on Risk, Insurance, and Economic Security*, ed. Georges Dionne and Claire Laberge-Nadeau, 151–173. Boston, MA:Springer US.

- Boyer, M. Martin.** 2000. “Centralizing Insurance Fraud Investigation.” *The Geneva Papers on Risk and Insurance Theory*, 25(2): 159–178.
- Braverman, Mark, Jieming Mao, Jon Schneider, and S. Matthew Weinberg.** 2017. “Multi-Armed Bandit Problems with Strategic Arms.” *arXiv:1706.09060 [cs, stat]*.
- Brockett, Patrick L., Richard A. Derrig, Linda L. Golden, Arnold Levine, and Mark Alpert.** 2002. “Fraud Classification Using Principal Component Analysis of RIDITs.” *Journal of Risk and Insurance*, 69(3): 341–371.
- Cameron, A. C., P. K. Trivedi, Frank Milne, and J. Piggott.** 1988. “A Microeconomic Model of the Demand for Health Care and Health Insurance in Australia.” *The Review of Economic Studies*, 55(1): 85–106.
- Cattaneo, Matias D.** 2010. “Efficient Semiparametric Estimation of Multi-Valued Treatment Effects under Ignorability.” *Journal of Econometrics*, 155(2): 138–154.
- Cattaneo, Matias D.** 2018. “Multi-Valued Treatment Effects.” In *The New Palgrave Dictionary of Economics*. 9208–9210. London:Palgrave Macmillan UK.
- Cornwell, Christopher, and William N. Trumbull.** 1994. “Estimating the Economic Model of Crime with Panel Data.” *The Review of Economics and Statistics*, 76(2): 360–366.
- Craig, Steven G.** 1987. “The Deterrent Impact of Police: An Examination of a Locally Provided Public Service.” *Journal of Urban Economics*, 21(3): 298–311.
- Crocker, Keith J., and John Morgan.** 1998. “Is Honesty the Best Policy? Curtailing Insurance Fraud through Optimal Incentive Contracts.” *Journal of Political Economy*, 106(2): 355–375.
- de Chaisemartin, Clément, and Xavier D’Haultfoeuille.** 2019. “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects.” National Bureau of Economic Research Working Paper 25904.
- Derrig, Richard A.** 2002. “Insurance Fraud.” *Journal of Risk and Insurance*, 69(3): 271–287.
- Dionne, Georges, and Kili C. Wang.** 2013. “Does Insurance Fraud in Automobile Theft Insurance Fluctuate with the Business Cycle?” *Journal of Risk and Uncertainty*, 47(1): 67–92.
- Dionne, Georges, Florence Giuliano, and Pierre Picard.** 2008. “Optimal Auditing with Scoring (Annex).” *Management Science*, 55(1): 58–70.

- Easley, David, and Nicholas M. Kiefer.** 1988. "Controlling a Stochastic Process with Unknown Parameters." *Econometrica*, 56(5): 1045–1064.
- Ehrlich, Isaac.** 1973. "Participation in Illegitimate Activities: A Theoretical and Empirical Investigation." *Journal of Political Economy*, 81(3): 521–565.
- Ferguson, Thomas.** 2004. *Optimal Stopping and Applications*. available at <https://www.math.ucla.edu/~tom/Stopping/Contents.html>.
- Gale, Douglas, and Martin Hellwig.** 1985. "Incentive-Compatible Debt Contracts: The One-Period Problem." *The Review of Economic Studies*, 52(4): 647–663.
- Gollier, Christian.** 1987. "Pareto-Optimal Risk Sharing with Fixed Costs per Claim." *Scandinavian Actuarial Journal*, 1987(1-2): 62–73.
- Goodman-Bacon, Andrew.** 2018. "Difference-in-Differences with Variation in Treatment Timing." National Bureau of Economic Research Working Paper 25018.
- Gossner, Olivier, and Tristan Tomala.** 2008. "Entropy Bounds on Bayesian Learning." *Journal of Mathematical Economics*, 44(1): 24–32.
- Graetz, Michael J., Jennifer F. Reinganum, and Louis L. Wilde.** 1986. "The Tax Compliance Game: Toward an Interactive Theory of Law Enforcement." *Journal of Law, Economics, & Organization*, 2(1): 1–32.
- Greenberg, Joseph.** 1984. "Avoiding Tax Avoidance: A (Repeated) Game-Theoretic Approach." *Journal of Economic Theory*, 32(1): 1–13.
- Greene, William.** 2004. "The Behaviour of the Maximum Likelihood Estimator of Limited Dependent Variable Models in the Presence of Fixed Effects." *The Econometrics Journal*, 7(1): 98–119.
- Guimarães, Paulo.** 2008. "The Fixed Effects Negative Binomial Model Revisited." *Economics Letters*, 99(1): 63–66.
- Hausman, Jerry, Bronwyn H. Hall, and Zvi Griliches.** 1984. "Econometric Models for Count Data with an Application to the Patents-R & D Relationship." *Econometrica*, 52(4): 909–938.

- Hörner, Johannes, and Andrzej Skrzypacz.** 2017. “Learning, Experimentation, and Information Design.” In *Advances in Economics and Econometrics.*, ed. Bo Honore, Ariel Pakes, Monika Piazzesi and Larry Samuelson, 63–98. Cambridge:Cambridge University Press.
- Huberman, Gur, David Mayers, and Clifford W. Smith.** 1983. “Optimal Insurance Policy Indemnity Schedules.” *The Bell Journal of Economics*, 14(2): 415–426.
- Keller, Godfrey, Sven Rady, and Martin Cripps.** 2005. “Strategic Experimentation with Exponential Bandits.” *Econometrica*, 73(1): 39–68.
- Khalil, Fahad.** 1997. “Auditing without Commitment.” *The RAND Journal of Economics*, 28(4): 629–640.
- Kirchler, Erich.** 2007. *The Economic Psychology of Tax Behaviour.* Cambridge University Press.
- Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez.** 2011. “Unwilling or Unable to Cheat? Evidence From a Tax Audit Experiment in Denmark.” *Econometrica*, 79(3): 651–692.
- Krawczyk, Michal.** 2009. “The Role of Repetition and Observability in Deterring Insurance Fraud.” *The Geneva Risk and Insurance Review*, 34(1): 74–87.
- Lancaster, Tony.** 2000. “The Incidental Parameter Problem since 1948.” *Journal of Econometrics*, 95(2): 391–413.
- Lehrer, Ehud, and Rann Smorodinsky.** 1996. *Merging and Learning.* Institute of Mathematical Statistics.
- Lehrer, Ehud, and Rann Smorodinsky.** 2000. “Relative Entropy in Sequential Decision Problems.” *Journal of Mathematical Economics*, 33(4): 425–439.
- Major, John A., and Dan R. Riedinger.** 2002. “EFD: A Hybrid Knowledge/Statistical-Based System for the Detection of Fraud.” *Journal of Risk and Insurance*, 69(3): 309–324.
- Meghir, Costas, and Jean-Marc Robin.** 1992. “Frequency of Purchase and the Estimation of Demand Systems.” *Journal of Econometrics*, 53(1): 53–85.
- Mittone, Luigi.** 2006. “Dynamic Behaviour in Tax Evasion: An Experimental Approach.” *The Journal of Socio-Economics*, 35(5): 813–835.

- Mookherjee, Dilip, and Ivan Png.** 1989. “Optimal Auditing, Insurance, and Redistribution.” *The Quarterly Journal of Economics*, 104(2): 399–415.
- Neyman, Abraham, and Daijiro Okada.** 1999. “Strategic Entropy and Complexity in Repeated Games.” *Games and Economic Behavior*, 29(1): 191–223.
- Picard, Pierre.** 1996. “Auditing Claims in the Insurance Market with Fraud: The Credibility Issue.” *Journal of Public Economics*, 63(1): 27–56.
- Picard, Pierre.** 2000. “On the Design of Optimal Insurance Policies Under Manipulation of Audit Cost.” *International Economic Review*, 41(4): 1049–1071.
- Picard, Pierre.** 2013. “Economic Analysis of Insurance Fraud.” In *Handbook of Insurance*. , ed. Georges Dionne, 349–395. Springer New York.
- Picard, Pierre, and Kili Wang.** 2015. “Insurance Fraud Through Collusion Between Policyholders and Car Dealers: Theory and Empirical Evidence.” working paper or preprint.
- Pinquet, Jean, Mercedes Ayuso, and Montserrat Guillén.** 2007. “Selection Bias and Auditing Policies for Insurance Claims.” *Journal of Risk and Insurance*, 74(2): 425–440.
- Praag, B. M. S. Van, and E. M. Vermeulen.** 1993. “A Count-Amount Model with Endogenous Recording of Observations.” *Journal of Applied Econometrics*, 8(4): 383–395.
- Reinganum, Jennifer F., and Louis L. Wilde.** 1986. “Equilibrium Verification and Reporting Policies in a Model of Tax Compliance.” *International Economic Review*, 27(3): 739–760.
- Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille.** 2007. “Social Learning in One-Arm Bandit Problems.” *Econometrica*, 75(6): 1591–1611.
- Rubinstein, Ariel.** 1979. “An Optimal Conviction Policy for Offenses That May Have Been Committed by Accident.” In *Applied Game Theory: Proceedings of a Conference at the Institute for Advanced Studies, Vienna, June 13–16, 1978*. , ed. S. J. Brams, A. Schotter and G. Schwödiauer, 406–413. Heidelberg:Physica-Verlag HD.
- Schiller, Jörg.** 2006. “The Impact of Insurance Fraud Detection Systems.” *Journal of Risk and Insurance*, 73(3): 421–438.
- Sutton, Richard S., and Andrew G. Barto.** 1998. *Reinforcement Learning: An Introduction*. MIT Press.

- Tennyson, Sharon.** 2008. “Moral, Social, and Economic Dimensions of Insurance Claims Fraud.” *Social Research*, 75(4): 1181–1204.
- Tennyson, Sharon, and Pau Salsas-Forn.** 2002. “Claims Auditing in Automobile Insurance: Fraud Detection and Deterrence Objectives.” *Journal of Risk and Insurance*, 69(3): 289–308.
- Townsend, Robert M.** 1979. “Optimal Contracts and Competitive Markets with Costly State Verification.” *Journal of Economic Theory*, 21(2): 265–293.
- Trevien, Corentin.** 2015. “Le Commerce de Détail d’optique Depuis 2003 - Insee Première - 1576.” <https://www.insee.fr/fr/statistiques/1560279>.
- Trumbull, William N.** 1989. “Estimations of the Economic Model of Crime Using Aggregate and Individual Level Data.” *Southern Economic Journal*, 56(2): 423–439.
- Varas, Felipe, Ivan Marinovic, and Andrzej Skrzypacz.** 2018. “Random Inspections and Periodic Reviews: Optimal Dynamic Monitoring.” Social Science Research Network SSRN Scholarly Paper ID 2905756, Rochester, NY.
- Viaene, Stijn, Richard A. Derrig, Bart Baesens, and Guido Dedene.** 2002. “A Comparison of State-of-the-Art Classification Techniques for Expert Automobile Insurance Claim Fraud Detection.” *Journal of Risk and Insurance*, 69(3): 373–421.
- Wald, A.** 1945. “Sequential Tests of Statistical Hypotheses.” *The Annals of Mathematical Statistics*, 16(2): 117–186.
- Whittle, P.** 1988. “Restless Bandits: Activity Allocation in a Changing World.” *Journal of Applied Probability*, 25: 287–298.
- Wooldridge, Jeffrey M.** 1997. “Multiplicative Panel Data Models without the Strict Exogeneity Assumption.” *Econometric Theory*, 13(5): 667–678.
- Wooldridge, Jeffrey M.** 1999. “Distribution-Free Estimation of Some Nonlinear Panel Data Models.” *Journal of Econometrics*, 90(1): 77–97.
- Wooldridge, Jeffrey M.** 2002. *Econometric Analysis of Cross Section and Panel Data*. MIT Press.
- Zhong, Weijie.** 2017. “Optimal Dynamic Information Acquisition.” *SSRN Electronic Journal*.

List of Tables

- 1.1 Definition of Variables 14
- 1.2 Summary Statistics 14
- 1.3 General auditing impact on number of submitted claims 22
- 1.4 Type-specific auditing impact on number of submitted claims 25

- 2.1 Period 1 audit decision x_1 as a function of period 0 beliefs and audit outcomes . 42
- 2.2 Definition of $C(\cdot)$ and $H(\cdot)$ 44
- 2.3 Definition of $C(\cdot)$ and $H(\cdot)$ 55

- 3.1 Marginal effects of auditing 74

- 4.1 Possible scenarios at a given period 109

List of Figures

- 1 Adaptation of defrauders and change in drug prescribers (Source CNAM²) 2
- 1.1 Timing of claim submission and auditing 12
- 1.2 Distribution of Claims 15
- 1.3 Claim submission and time dependence 16
- 1.4 Directed Acyclic Graph for equation (1.4) 19
- 1.5 Counterfactual Submissions at the individual level 23
- 1.6 Deterrence as a probability point mass shift to the left 24
- 2.1 Updating Priors According to the Auditing Status ($p_D = 0.9, p_H = 0.1$) 39
- 2.2 Updating functions for different parameters 40
- 2.3 Period 0 priors and period 1 auditing 42
- 2.4 Period 1 beliefs as a consequence of Period 0 audit outcomes 43
- 2.5 $K(\pi_0)$ for ($p_D = 0.9, p_H = 0.1, c = 0.5, \gamma = 1$) 45
- 2.6 Learning vs Myopic 46
- 2.7 $x_0(\pi_0)K(\pi_0)$ for ($p_D = 0.9, p_H = 0.1, c = 0.5, \gamma = 1$) 46
- 2.8 Auditing thresholds $\pi^*(\ell)$ and $\pi^{**}(\ell)$ with a variable claim value ℓ 48
- 3.1 Updating Functions 68
- 3.2 Simulated convergence of priors to the true beliefs 69
- 3.3 Distribution of priors after repeated audits ($T = 1000$). 70
- 3.4 Myopic objective, policy and value functions. 71
- 3.5 Time t policy, objective and value functions 73
- 3.6 Levels of Exploration/Exploitation 74
- 3.7 Finite Horizon Thresholds and Value Function 75
- 3.8 Infinite Horizon Thresholds and Value Functions 77
- 3.9 Optimal auditing as an optimal stopping rule 79

3.10	6-armed bandit with a capacity constrained audit ($k = 2$)	82
3.11	Updating functions for different parameters	85
3.12	Optimal thresholds depend on β	87
3.13	Polarized and extreme fraud behaviors induce more exploration	87
4.1	Instantaneous Equilibrium Auditing and Fraud	111
4.2	Instantaneous Equilibrium Cost	112
4.3	Updating functions for $(p_0 = 0.2, p_H = 0.1)$	114
4.4	Audit, fraud and average discounted cost at $(t, T) = (1, 2)$	117
4.5	Audit, fraud and average discounted cost at $(t, T) = (1, 3)$	120
4.6	Multi-period costs in the vicinity of $\hat{\pi}$	121
4.7	$\bar{C}_1(x_1, y_1, \pi)$ for $T = 3$	122
4.8	Belief trajectories for a type D under myopic strategies	123
4.9	Belief trajectories for a type H under myopic strategies	124
4.10	Proof cases	127

Titre : Effets de dissuasion et d'apprentissage dans l'audit de la fraude à l'assurance

Mots clés : Fraude, assurance, dissuasion, apprentissage, interactions répétées

Résumé : La fraude à l'assurance est une menace majeure pour les marchés assurantiels dont le traitement passe par la conception de politiques d'audit crédibles et ciblées. Cette thèse étudie les mécanismes de dissuasion et d'apprentissage dans l'audit de la fraude à l'assurance, notamment quand des prestataires (garagistes, opticiens, etc.) jouent le rôle d'intermédiaires entre l'assureur et les assurés. Le premier chapitre est une étude empirique des effets dissuasifs de l'audit, fruit d'une collaboration avec IBM France et PRO BTP, dans le cadre du déploiement de la solution anti-fraude Solon. Cette analyse montre que les audits subis par un opticien ont pour effet de réduire sa fraude future. Plus spécifiquement, l'effet dissuasif est d'autant plus fort que la menace d'audit est crédible, soulignant l'importance de l'engagement dans la lutte contre la fraude. Les deuxième et troisième chapitres étudient un problème d'audit dynamique où l'information joue un rôle central. L'auditeur y interagit de façon répétée avec des prestataires non-stratégiques et peut apprendre à propos de

leur propension à frauder sur la base des résultats de l'audit. Le deuxième chapitre utilise un modèle à deux périodes pour mettre en évidence cet effet d'apprentissage, dont la conséquence est qu'il est optimal d'auditer plus intensément au début de la relation. Le troisième chapitre étend ce modèle à un nombre arbitraire ou infini de périodes, et montre que l'audit optimal est d'autant plus intense qu'il reste un nombre important d'interactions futures. L'intuition réside dans le fait que davantage d'audit au présent, quoique coûteux, a des répercussions informationnelles positives sur toutes les périodes suivantes. Enfin, le quatrième chapitre réunit les mécanismes de dissuasion et d'apprentissage dans un même modèle dynamique de réputation, avec des prestataires stratégiques. Il montre l'existence d'une forme de dissuasion réputationnelle où l'apprentissage transforme la dissuasion en une menace intertemporelle. En d'autres termes, un prestataire sera davantage dissuadé dans le présent car il risque de détériorer sa réputation future s'il se fait attraper en train de frauder.

Title : Deterrence and learning effects in insurance fraud audits

Keywords : Fraud, insurance, deterrence, learning, repeated interactions

Abstract : Insurance fraud is a serious threat to insurance markets and is tackled through the design of credible and targeted auditing policies. This thesis studies the deterrence and learning mechanisms of insurance fraud audits, especially when service providers (car repairers, opticians, etc.) act as intermediaries between the insurer and the policyholders. The first chapter is an empirical assessment of the deterrence effects of auditing. It was conducted in collaboration with IBM France and PRO BTP, in the context of the deployment of the Solon counter-fraud solution. This assessment shows that incurred audits decrease an optician's subsequent fraud. More specifically, the more credible the audit threat, the stronger this deterrence effect, emphasizing the importance of commitment in counter-fraud efforts. The second and third chapters examine a dynamic auditing problem where information plays a central role. The auditor interacts repeatedly with non-strategic service providers and can learn about their propensity to defraud

from the auditing outcomes. The second chapter relies on a two-period model to show the existence of this learning effect, whose consequence is that it is optimal to audit more at the beginning of the relationship. The third chapter extends this model to an arbitrary or infinite number of periods, and shows that the further away the time horizon, the larger the optimal auditing efforts. Intuition stems from the fact that more auditing in the present, though costly, has a positive informational impact on all future periods. Finally, the fourth chapter combines the deterrence and learning mechanisms in the same dynamic reputation model, with strategic service providers. It reveals a reputation-based deterrence effect, where learning turns deterrence into an intertemporal threat. In other words, a service provider will be deterred more strongly in the present because of the risk of seeing his future reputation deteriorate if he gets caught defrauding.