

Game-Theoretic Analysis of Development Practices: Challenges and Opportunities

Carlos Gavidia-Calderon^{a,*}, Federica Sarro^a, Mark Harman^a, Earl T. Barr^a

^a*Department of Computer Science, University College London, London, UK*

Abstract

Developers continuously invent new practices, usually grounded in hard-won experience, not theory. Game theory studies cooperation and conflict; its use will speed the development of effective processes. A survey of game theory in software engineering finds highly idealised models that are rarely based on process data. This is because software processes are hard to analyse using traditional game theory since they generate huge game models. We are the first to show how to use game abstractions, developed in artificial intelligence, to produce tractable game-theoretic models of software practices. We present Game-Theoretic Process Improvement (GTPI), built on top of empirical game-theoretic analysis. Some teams fall into the habit of preferring “quick-and-dirty” code to slow-to-write, careful code, incurring technical debt. We showcase GTPI’s ability to diagnose and improve such a development process. Using GTPI, we discover a lightweight intervention that incentivises developers to write careful code: add a *single* code reviewer who needs to catch *only* 25% of kludges. This 25% accuracy is key; it means that a reviewer does not need to examine each commit in depth, making this process intervention cost-effective.

Keywords:

Game Theory, Empirical Analysis, Technical Debt, Software Engineering Practices

1. Introduction

*What’s important is not just to develop the technology;
it’s to develop the processes. —Hal Abelson*

Modern society revolves around software. We use it to communicate, understand global warming, operate machines, and decide what to buy. Like the human beings who write it, software is fallible. Software engineers rely heavily on tools to write and maintain robust software: tools that find bugs, repair bugs, or help developers avoid introducing them in the first place. Development tools are important, but so are the processes that use them. In the words of Linus Torvalds, “All the really stressful times for me have been about process: they haven’t been about code” [1]. Despite

*Corresponding author

Email addresses: carlos.gavidia.15@ucl.ac.uk (Carlos Gavidia-Calderon), f.sarro@ucl.ac.uk (Federica Sarro), mark.harman@ucl.ac.uk (Mark Harman), e.barr@ucl.ac.uk (Earl T. Barr)

their importance, researchers and practitioners have lacked the tooling to build bespoke formal, yet still intuitive and explainable, models for software processes. Instead, they had to rely on generic, coarse-grained models or models difficult to elicit and often hard to understand.

Developers, customers, and managers cooperate and compete to write software. These interactions define a software process. Examples include effort estimation, where managers underestimate to win a project bid, while technical leaders overestimate to minimise risk [2, 3], and defect prioritisation, where end users exaggerate the importance of their bugs to obtain their fix quickly, while the engineering team needs accurate priorities (business value, not technical severity) to maximise the value delivered per release [4].

These software processes can be viewed as games. Game theory studies mathematical models of conflict and cooperation [5], and has already been used in the analysis of complex scenarios like arms races between nations or duopolies in markets [6]. One of its more important findings is the *Nash equilibrium* (NE): a game outcome where players have no incentive to deviate by choosing different actions. Real-world scenarios where the players have mastered the game — like professional sports — converge on their NE [7, 8]. When we model software processes as games, we can compute their NE to diagnose their problems and prescribe fixes when needed.

We believe that most software processes are emergent phenomena that arise to solve problems. Since they are not designed, many are suboptimal: they misalign a player’s payoffs with the overall goals of a process, permitting, even encouraging, players to act in ways that undermine collective goals. These games have an undesirable NE. *Mechanism design* identifies and fixes games whose NE is undesirable. Players use strategies to decide their actions. *Game designers* apply mechanism design to formulate games whose NE favours desirable strategies. We are proposing that process engineers [9] be game designers, who craft software processes considering cooperation incentives and conflict avoidance mechanisms.

Despite game theory’s obvious applicability to software processes, our survey of the state of the art (Section 3) shows that software engineering researchers have not fully exploited game theory, limiting their analysis to idealised models. The reason is that software processes are inherently temporal: they are not one-off events; they comprise interactions over time. Game theory models these scenarios with extensive form games. This formalism is intricate and its analysis must cope with exponentially large game trees. Even constrained versions of poker have trees with 10^{18} leaves [10].

Our research agenda exploits abstraction to build tractable game-theoretic models of real-world software development scenarios. *Empirical game-theoretic analysis* (EGTA) is an instance of a large set of abstractions for dealing with extensive form games [11]. EGTA converts an extended form game into a normal-form game that only considers a subset of the actions available in the full game and the payoffs are expected values obtained through simulation. EGTA relies heavily on data for simulation input analysis and output validation. Fortunately, software engineers can now easily access a wealth of data building and using software products [12]. Our proposed solution — called Game-Theoretic Process Improvement (GTPI) — relies on this data and EGTA models to improve software development processes (Section 4).

Many tasks can be partially completed. At one end of the scale, one can use a quick and dirty short term fix; at the other, one takes the time and care to find and implement a complete and lasting fix. Neither approach is always better. Quick code can help a development team beat a competitor

to market, but, taken too far, turn a codebase into an unmaintainable mess. When developers are not racing, we prefer more deeply considered code. Some development workflows settle into a suboptimal point along this continuum. Lavallée and Robillard observed that some software teams had a tendency to prefer quick-cheap solutions over time-consuming ones, incurring technical debt due to fear of exceeding their budget [13]. We use this problem to showcase GTPI.

Our contributions follow:

- A survey on game-theoretic analysis of software development practices.
- An introduction to EGTA for software engineers; and
- A complete, end-to-end analysis and fix of budget-driven technical debt: we diagnose it using game theory, then apply mechanism design to suggest a simple and inexpensive change — the addition of a single reviewer — and validate the new process via its NE.

2. Background

In this section, we review key concepts from game theory using a software engineering scenario to illustrate them. We also discuss the issue of game representation: the number of players, actions and rounds over time have a large impact on the size of the game-theoretic model.

In game-theoretic context, a *game* is a scenario where multiple self-interested agents, or *players*, interact and their actions impact the perceived utility of all of them. This definition fits nicely with the traditional definition of game. For example, in rock-paper-scissors winning is a function of your play — selecting an *action* — and the play of your opponent.

Consider a software project that hires two freelance engineers: a frontend engineer and a backend engineer. Their contract stipulates \$50 upfront and an additional \$50 upon completion. If the project goes live before the deadline, they each receive an additional \$50. They have two actions: to cooperate and work together or to work individually, not cooperate and ignore each other. If they cooperate, they will certainly finish before the deadline. Otherwise, they will certainly finish their individual tasks, but lack of integration means the project will not go live. If only one cooperates, the isolationist will finish their task quickly and free themselves to take another contract, while the cooperating freelancer will not complete their task. We represent this game with a *payoff table* — shown in [Table 1](#) — that contains payoff information per player given the actions they performed.

A *strategy* defines the behaviour of the player in a game. Strategies fall into two categories. In *pure strategies*, players select a single action; in *mixed strategies* players randomise over actions according to a probability distribution [14]. In an NE, all the players adopt the best response to their opponents' strategies: any deviation lowers a player's payoff. This explains the stability of an NE. Nash proved that every finite game has at least one NE [15].

To put this in context, we return to our freelancers. When one freelancer cooperates and the other does not, the outcome is unstable since only the uncooperative player has adopted the best response, while the cooperative freelancer would be better off not cooperating. The NE for both players is to adopt the same strategy: ignore each other. Any deviation from this strategy allows their colleague to take advantage. Hence, NE analysis predicts the freelancers work independently and obtain \$100 each without completing the project. However, in the unstable scenario where they

Table 1: Pay-off matrix for the freelancer’s dilemma: each cell contains the payoff in dollars obtained by Freelancer A and Freelancer B given their choice to cooperate or not.

	<i>B: cooperate</i>	<i>B: not cooperate</i>
<i>A: cooperate</i>	A = \$150, B = \$150	A = \$50, B = \$200
<i>A: not cooperate</i>	A = \$200, B = \$50	A = \$100, B = \$100

work together, they obtain \$150 and a satisfied client with a completed project. This dilemma faced by our freelancers is an instance of the *prisoner’s dilemma*: a game used to explain cooperation — or the lack of it — in real-world scenarios like global warming and cigarette advertisement.

Game-theoretic models rely on strong assumptions regarding player knowledge. For NE convergence, game structure, player rationality and player beliefs must be common knowledge. These assumptions do not always align with human behaviour [16]. In contrast, *bounded rationality* approaches model rationality within the limits of attention, information, and mental capabilities of the decision maker [17]. *Prospect theory* is a bounded rationality model that describes decision-making under uncertainty. According to it, agents make decisions based on utility relative to current wealth — gains or losses— rather than the absolute wealth obtained. Prospect theory has been successfully applied to the design of drone delivery systems [18] and to the study of smart-grid adoption [19]. *Cognitive hierarchy* also relaxes game-theory’s rationality assumptions. Each player in a cognitive hierarchy model has a level, which determines the number of strategic reasoning iterations the player can perform. Level-0 players choose actions uniformly at random; players of superior levels are “smarter” in their decision-making. *Psychological games* extend the utility function definition to also depend on the players beliefs and their beliefs regarding the other players. This includes social norms and emotions in the game-theoretic model. In this work, we adopt the conventional game-theory approach that predicts convergence to NE under its rationality assumptions. We believe teams of experienced developers working over well-know codebases, like in open-source projects or in mature software maintenance teams, meet these assumptions. Our results are already strong despite being conservative and imprecise and they can only improve by adopting a bounded rationality approach closer to team behaviour. This is left as future work.

Game theory has been widely applied. It has already been used to optimise student recruitment in American hospitals, improve the auction of telecom operating licences in Britain and to analyse transport networks in New York [20]. In *static games*, players select their actions independently and simultaneously, like in rock-paper-scissors. *Dynamic games*, like chess or poker, unfold over time and require players to sequence their actions. The most common representation for dynamic games is the *extensive form* [16]. Extensive-form games can be described via game trees, like the one in Figure 1. The non-terminal nodes of the game tree are called *choice nodes*. In each choice node, a player is required to select an action and the available actions correspond to the outgoing edges of the choice node. A *pure strategy in an extensive form game* defines which action to take at each player’s choice node, while a *mixed strategy* is a probability distribution over these pure strategies [14]. The size of the game tree can become inconveniently large when dealing with real-world scenarios. These sizes quickly become unmanageable for NE calculation algorithms, hence abstraction is needed.

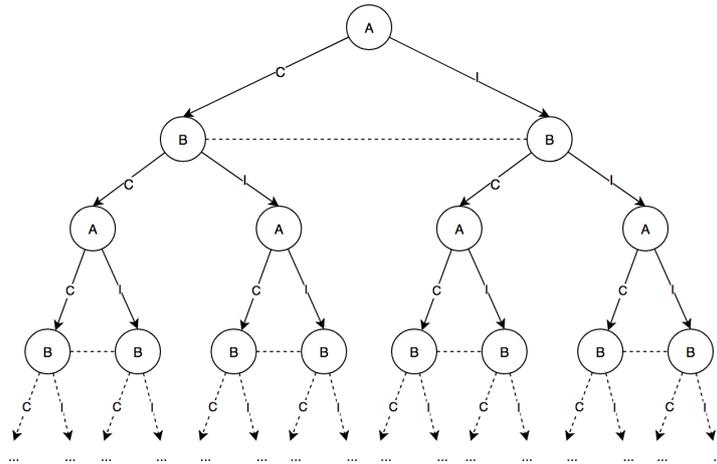


Figure 1: Extensive form game for a multi-project freelancer’s dilemma: this figure contains the part of the tree corresponding to the first two projects. The size of this representation grows with the number of projects, freelancers and actions available to them. EGTA (Section 4) is crucial for managing this explosive growth.

3. A Survey of Game Theory Applied to Software Engineering

In the Section 2, we used game theory to analyse freelancers in a software project and showed its potential as a formal technique to analyse development practices. Other popular approaches are less formal but based on the consultants’ hard-won experience, although they sometimes lack arguments to back their adoption [21]. Bayesian networks have also been proposed for process improvement, and due to their use of models and their formal approach they are the closest to game theory [22]. We start this section by exploring Bayesian networks, then review the work published so far in game-theoretic applications to software engineering.

Bayesian Networks for Process Improvement. Scenarios like the tragedy of the test suite (Section 5) can be analysed with tools from decision theory, like the Bayesian network in Figure 2, but this approach fails to capture strategic interactions *between agents*. Decision theory can be seen as “the study of games against nature”, where nature behaves stochastically instead of been driven by self-interest [23]. Game theory models agents’ interests and how their interactions affect these interests. In game theory, players, actions, and strategy profiles are all first-class elements. Model elicitation is the coal face of process improvement, and it is here that game theoretic models shine, because their primitives are easy to understand and are a natural fit for modelling software processes. In contrast, building Bayesian networks rests on eliciting random variables that can have unclear conceptual boundaries. Bayesian networks can, of course, model game theoretic models, at the cost of further complicating their models.

Game Theory for Software Engineering. Game theory has been used before for exploring conflict and cooperation in a software development context. Below, we survey game-theoretic analysis of five types of interactions: interactions between developers, development and testing conflicts, interactions driven by management, role-neutral interactions, and conflicts external to the development

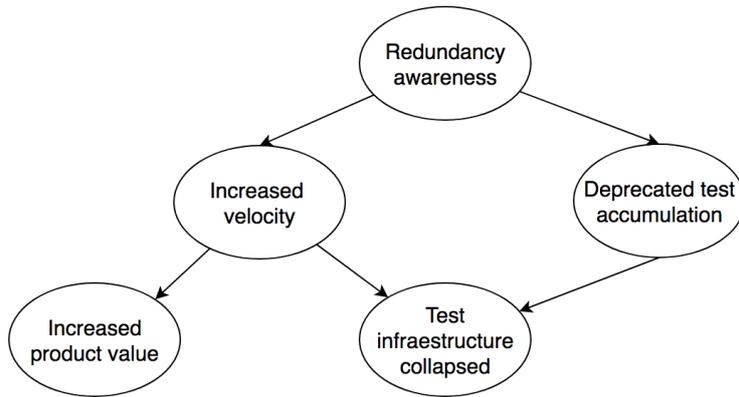


Figure 2: A Bayesian network for the tragedy of the test suite (Section 5): redundancy awareness, understood as the perception that underperforming can get a developer fired, influences the increase in velocity and the accumulation of deprecated tests. Higher velocity and the accumulation of deprecated tests influence the collapse of testing infrastructure.

team. However, the proposed models do not scale to global software development scenarios with large teams and fast releases [11].

Researchers have modelled strategic scenarios between developers on a software team. Lagesse explored task assignment between developers, using cooperative game theory to obtain an ideal assignment considering the developers’ preferences [24]. Hasnain *et al.* modelled stand up meetings, where developers had the option to work or shirk responsibility [25]. Bavota *et al.* approached refactoring as a game between developers, where developers compete for refactoring opportunities on a common class to refactor [26]. Kitagawa *et al.* explored the collaboration between code reviewers inside a software team, proposing that this scenario constitutes an instance of the snowdrift game [27]. These papers succeed in identifying conflicting goals between software developers. In the same fashion, all the games presented in this paper — the freelancer’s dilemma (Section 2), technical debt (Section 4), and the test suite tragedy (Section 5) — fall in this category. However, since they use techniques from classical game theory, their models cannot scale to large software teams operating over time, as explained in Section 2.

Other authors have explored scenarios where developers are in conflict with the testing team. Feijs models testing as a game between developers and testers, where the actions for both of them are to do a good job or a poor job [28]. Kukreja *et al.* modelled testing as a security game: testers are defenders that have a limited budget for test case execution while developers are the attackers who commit poor quality code [29]. The tension between the testing and the developer role fits naturally into the game-theoretic domain, and we expect to study it in the future. As with previous work, they rely on classic game theoretic approaches that limit their scalability.

The connection between software project management and game theory was identified by Boehm and Ross [30]. Models of software process improvement normally have the project manager as player. Hata *et al.* modelled collaboration in an open-source project as a game between project leaders and contributors; the project leader’s goal is to facilitate contribution and his actions include doing nothing, writing documentation, or writing installation scripts [31]. Yilmaz *et al.* focused on team members’ characteristics and personalty traits and their impact on productivity [32, 33]. Bacon

et al. proposed viewing software development as an economy. They present software estimation as an example, where project managers are rewarded by accurate estimations and developers benefit from quick feature delivery [34]. Following that vision, Rao *et al.* suggested a software maintenance game where developers compete by submitting deep or shallow fixes, comparable with the fix-kludge scenario we presented in Section 4 [35]. Rao *et al.* also mention the problem of scaling game-theoretic models. They address this by using mean field games, which support a large number of players. Some interesting topics covered by these authors are mechanism design [32, 34, 35], data-based validation [31] and game abstraction [34]. However, they studied each of these topics in isolation. A key contribution of this work is to integrate all these concepts into a formal framework for process improvement at scale, from inception to deployment.

There are also approaches that model team members regardless of their role in the software process. Wang and Redmiles [36] explore the impact of informal talk in the development of trust using evolutionary game theory. Hassan and Dubinsky [37] described how cooperation and defection would impact specific software development practices. Again, since no game abstraction approach was adopted, those game-theoretic models do not scale.

Strategic interactions outside the software team have also been explored, especially with end-users and clients. Since these proposals focus on clients and users, they obviate software development practices, which is a focus of this research agenda. García-Galán *et al.* reconcile conflicting user configurations in product design by adopting a cooperative game theory approach [38]. Grechanik and Perry model software development as a game with three types of players: managers, developers and customers [39]. Oza explores that idea that software outsourcing can be modelled as a game between software vendors and their clients, configuring a prisoner’s dilemma scenario [40]. Sazawal and Sudan model software evolution as an extensive form game between the software designer and its end-users [41]. Kumar *et al.* proposed a novel marketplace style for cloud computing providers and customers based on the double auction mechanism [42].

4. Research Agenda: Game-Theoretic Software Process Improvement

Since software development is inherently temporal, time is a critical dimension of most software processes. For example, the Scrum process framework divides development in time boxes, called *sprints*, where development tasks are assigned and prioritised. At the end of each sprint, the status of the project is assessed and the project plan is refined accordingly. Other process frameworks, like the Unified Process, also propose guidelines for sequencing tasks and organising them over time. Dynamic games and the extensive-form (Section 2) are more suitable for representing software processes than static games, given static games inability to handle time [16]. However, real-world software development — with large teams working over multiple releases — generate immense game trees. Hence, abstraction is needed. *Empirical game-theoretic analysis (EGTA)* is one of these abstraction approaches [43] [44]. EGTA reduces the action space by restricting it to a set of *heuristic strategies*, which are a subset of strategies that are of interest to the game designer. In our freelancer scenario, let us assume that the freelancers work together for an undetermined number of projects under the contractual conditions of Table 1: strategies can go from “not cooperate on any project” to “imitate my colleague’s last action”. In this strategy space, the game designer needs

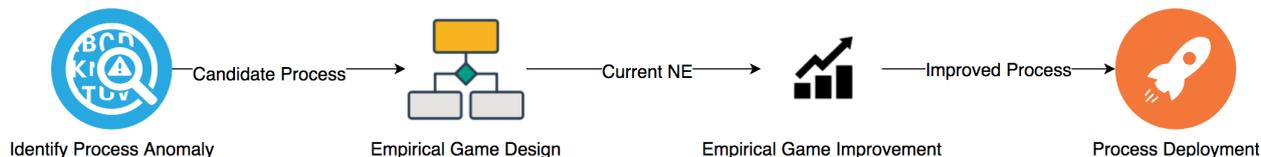


Figure 3: The GTPI approach starts by detecting a candidate process that exhibits an aberrant behaviour. Such a process is then modelled using EGTA and the Nash equilibrium is obtained. In case the equilibrium is not the one desired, the EGTA model is updated iteratively until a desirable one is obtained. Finally, the improved process is adopted by the team.

to select a subset of these strategies. Heuristic strategies can be obtained from data or hand-crafted following first principles, based on the properties of the system under study [44].

EGTA uses a simulation-based *heuristic payoff table* instead of the game tree of the extended form representation. This table has an entry for each action combination, where the actions available for each player are heuristic strategies. Each entry also contains the expected payoff for each player given the actions they perform, which are obtained through simulation. The heuristic payoff table can then be processed by a game solver to obtain its NE.

Game-theoretic Process Improvement (GTPI) is our proposed process improvement framework. It relies on EGTA abstractions to model software processes, diagnose process anomalies and prescribe solutions to remove them. It is summarised in Figure 3. Anomalous processes are often itchy: something about them mystifies or annoys their participants. Some behaviours that are apparently irrational appear due to incorrect incentives. *Identifying process anomalies* is the first step of our approach.

Workers of every discipline, when approaching a task, need to choose between “cutting corners” or not. Technical debt is a software development instance of this seminal problem [13]. Practitioners [45, 46] and researchers [13, 35] have identified that software teams have an incentive to deliver sub-optimal but quick solutions — like kludges — instead of optimal but time-consuming ones, like proper fixes. One of the causes is the pressure put on teams to deliver on time and under budget, which triggers them to maximise the number of features delivered regardless of quality. Individual developers too balance getting things done quickly versus getting them done right. Developers face this dilemma repeatedly and, of course, they seek a Goldilocks solution¹.

The next step is the *empirical game design*, where we model the process to improve. Underlying Figure 4 are two models. One is a simulation model designed to capture real world behaviour. To the simulation model, we apply heuristic strategies to produce the other, an empirical game-theoretic model in the form of pay-off table. Simulation of software processes is a well-developed area, with plenty of options and paradigms [9]. Here, we adopted discrete-event simulation because it is easy to reproduce and evaluate [47].

The parameters for our simulation model are described in Table 2. Technical debt can arise because developers write sloppy code under pressure. Why? Game theory suggests the answer is incorrect incentives, so the first modelling question to ask is “How are developers rewarded?”. The answer is not directly money or reputation, so we look to the output of the process; F_i features

¹“Goldilocks and the Three Bears” is a 19th century fairy tale, in which a girl, given a series of three-way choices, consistently chooses the one between the extremes.

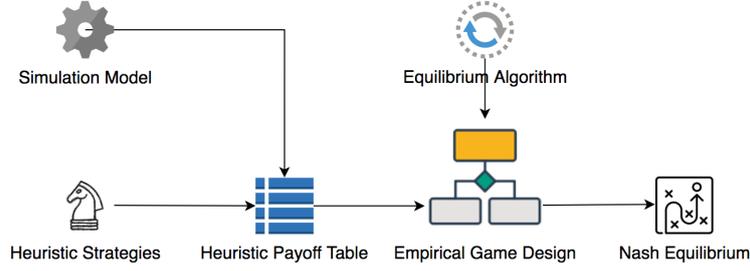


Figure 4: Empirical game design: the process engineer builds a heuristic payoff table using heuristic strategies as actions: The entries of this table are expected payoff values obtained via simulation. They can later use any algorithm to calculate the NE of the empirical game, once the payoff matrix has been built.

Table 2: Input and output variables of the simulation model of the development team: the output variable F_i corresponds to the payoff function of developers. The process engineer needs to work with the customer to identify the relevant variables of the process under analysis.

Input	Description
T	Resolution time for work items in days.
T_a	Resolution time for work items coded with action $a \in \{f, k\}$.
I	Work item arrival probability at the “To-Do” column per day.
S_i	Heuristic strategy adopted by developer i .
D	Number of developers available.
R	Rework a work item after it is placed in “Done”.
R_a	Rework a work item coded with action $a \in \{f, k\}$ after it is placed in “Done”.
N	Iteration duration in days.
Output	Description
F_i	Work items finished by developer i .

per release. We use F_i as the game’s payoff function below. F_i is governed by D , the number of developers and the release periods of N days. Kanban is popular among agile teams, so the development team in our model works a Kanban board with three columns: To-Do, In-Progress and Done². Tasks appear as To-Do’s with probability I . On average, a task project takes T days to reach “Done” and increase F_i .

Technical debt is a *tragedy of the commons*: the codebase quality, the shared resource or *common*, progressively degrades as the team pushes kludges. So, we define our model to make kludges faster to code than fixes at the expense of codebase quality and with a higher risk of rework. To this end, we model the probability of reworking a “Done” task as due to a bug. Fixes take T_f time and are R_f likely to require rework, while kludges take T_k time and are R_k likely to require rework. Kludges reduce the quality of the codebase and increase the average resolution time by Q_k . Of the myriad actions developers can take, our simulation model gives a developer only two:

²Kanban boards visualise workflows. They have several columns and work items flow between them [46].

write a kludge or code a time-consuming fix, as [Figure 5](#) shows. The simulation model must be validated against data with respect to the targeted behaviour, which is defined in cooperation with the customer. In essence, one statistically compares the simulation outputs with actual process data [\[48\]](#). Consolidating data across sources is non-trivial and time-consuming: we discuss how one might approach this problem in [Section 5](#).

The EGTA approach also requires a set of heuristic strategies. These strategies model player behaviour, and the game designer employs mechanism design to control their adoption in the empirical game improvement step. Ideal candidates are behaviours we want to encourage or discourage. To obtain the strategies players actually adopt, we extract data from bug tracking systems, source code repositories, eliciting them from the customer or domain experts, or other relevant data sources. When the available data is insufficient, we turn to experts to expand the strategy catalogue, keeping in mind that the number of strategies determines the size of the game.

Our model uses two strategies that we believe are worth exploring. The first commits proper fixes up to a week before the release, when, due to pressure, it shifts to pushing kludges. This behaviour favours fixes over kludges, most of the time. The second is sensitive to the work items accumulating in the “To-Do” column: when the “To-Do” backlog exceeds 2 items, it starts committing kludges. In a heavily loaded project, this strategy favours kludges. While actual player strategies will be some more subtle mix of fix and kludge, we have picked two extremal strategies to study how to reduce kludging. In focusing on extremal strategies, we are simply seeking to make our process resilient to unwanted behaviour, following a long tradition in mechanism design [\[49\]](#).

Our development model is simple by design in order to capture fundamental behaviour. Our key assumption is that the aberrant behaviour that we seek to capture has such strong signal that a simple model *can* capture it. Also, the explanatory power of simple models greatly favours process adoption; we believe customers and process performers will be more positively inclined towards process changes justified using concepts they can understand. Further, capturing the essence of the phenomena under study in a simple model generates more flexible models. Indeed, cutting corners is not unique to software. With just a few modifications, we could easily port our simulation model to non-software domains. Below, we show how our model suggests a simple, inexpensive solution to budget-driven technical debt: “All models are wrong but some are useful”, G. Box.

From the simulation model and the heuristic strategies, we obtain the heuristic payoff matrix by simulating each of the table entries and including the expected values over a number of iterations. This payoff matrix represents our game-theoretic model. Once we have the payoff matrix, we obtain its equilibria. Nash equilibrium calculation is a vast and diverse area with many options [\[50\]](#). In this paper, we rely on the Gambit game solver: a software tool that contains ready-to-use implementations of equilibrium calculation algorithms, facilitating quick experimentation [\[51\]](#).

An Eclipse Platform annual release has the simulation parameters $T = 29.79$ and $N = 360$ after analysing its Bugzilla data [\[52\]](#). Using our heuristic strategies and setting $D = 2, I = 1.0, T_f = 1.1 \times T, R_f = 0.9 \times R, T_k = 0.75 \times T, R_k = 1.05 \times R, Q_k = 0.05 \times T$, where $R = 0.069$ is raw rework probability in the Eclipse data. We have extracted these parameters to use real-world values. We want to emphasise that we do *not* assume that the Eclipse platform is kludge-prone. We simulated each heuristic payoff table entry for 100 releases and recorded the average F_i per developer. [Table 3](#) shows the resulting heuristic payoff matrix. Gambit finds a single equilibrium where both developers adopt the kludge-intensive strategy. This outcome matches the itchy behaviour identified at

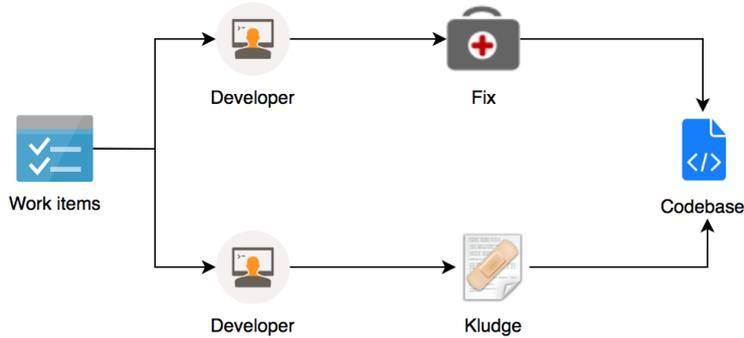


Figure 5: Developers in our technical debt simulation: the implementation of work items can be fixes or kludges. Fixes demand more time but are less likely to require rework. Kludges are quick but are more likely than fixes to be reworked. Also, kludges negatively impact codebase health.

Table 3: Payoff matrix after the empirical game design stage in Figure 3: it has a single Nash equilibrium where developer A and developer B adopt only the kludge-intensive heuristic strategy.

	<i>B: fix-intensive</i>	<i>B: kludge-intensive</i>
<i>A: fix-intensive</i>	$A = 9.80, B = 9.80$	$A = 8.34, B = 10.77$
<i>A: kludge-intensive</i>	$A = 10.77, B = 8.34$	$A = 9.06, B = 9.06$

the first stage of GTPI, and our empirical game model offers an explanation: developers have an incentive to kludge instead of generating proper fixes.

Following design, we proceed with the *empirical game improvement* step. Our goal is a process whose corresponding empirical game has a single equilibrium where the heuristic strategies we believe beneficial have high probability. To remove technical debt, we seek an equilibrium where the fix-intensive strategy has a significantly higher probability than the kludge-intensive one.

The analysis of Table 3 shows that the advantage of kludges — a reduced coding time — is worth their cost — a higher rework probability — so it becomes dominant at equilibrium. The absence of a code quality control mechanism before committing makes kludges very cheap, which translates in a progressive deterioration of codebase quality. Thus, we knew any solution should make kludges more expensive. Adopting pair programming can accomplish this but, due to its cost, we did not consider it. We posited that adding a part-time experienced code reviewer who can detect 10% of the kludges would be sufficient and also cost effective. So, we updated the simulation model with $R_k = 10\%$ and use it to produce payoff values for a new heuristic payoff table. In this configuration, Gambit found three equilibrium profiles and in one of them both agents perform the kludge-intensive strategy, which was far from ideal. However, when we set $R_k = 25\%$ and rebuild the payoff matrix, shown in Table 4, its equilibrium analysis produced the desired outcome: a single equilibrium where both players adopt the fix-intensive strategy.

The last step is the *process deployment*. The results of this analysis would be pointless without a feasible deployment strategy. Our improved process is deployable: most software teams already review commits. Since we only require a 25% kludge detection accuracy, the reviewer needs only to perform a lightweight quick-pass.

GTPI is a general approach to improving development processes. Here, we applied it to the

Table 4: Payoff matrix after the empirical game improvement stage in Figure 3: it has a single Nash equilibrium where developer A and developer B adopt only the fix-intensive heuristic strategy.

	<i>B: fix-intensive</i>	<i>B: kludge-intensive</i>
<i>A: fix-intensive</i>	$A = 9.79, B = 9.79$	$A = 8.40, B = 7.77$
<i>A: kludge-intensive</i>	$A = 7.77, B = 8.40$	$A = 6.83, B = 6.83$

technical debt problem. In other work, we built a tool, which we christened TaskAssessor, that applies GTPI to diagnose and remove priority inflation [53].

5. The Challenges of Applying GTPI

Adopting GTPI is challenging; data gathering, technical validation, and securing customer and performer acceptance is hard. We will use the following scenario to illustrate these challenges.

Imagine a large organisation whose flagship software product has been under constant development for years. Even though software products tend to grow by adding features, over time they also lose some of them. These lost features obsolete some tests. The problem is that this subset is largely unknown and, in practice, rarely removed. The main reason is that the engineering time needed to identify these useless tests with certainty is significant. Also, mistakenly removing a useful test can cost a software engineer their job. There are some — perhaps folkloric — stories of engineers losing their jobs because they removed a test that would have detected a security vulnerability. These factors explain why potentially useless tests are rarely removed, as illustrated in Figure 6. This wastes testing resources.

We call this scenario the *tragedy of the test suite* since it constitutes a tragedy of the commons: the test infrastructure is a shared resource that self-interested software engineers spoil by constantly adding tests without removing unneeded ones. Over time, maintaining and running the test infrastructure becomes more expensive, and test execution takes more time, which is harmful for the engineering team and the organisation as a whole. Adopting a test prioritisation technique can mitigate the problem but it does not in general remove it. Ideally, when software engineers modify, or remove, functionality that makes some tests redundant, they should proceed to remove them from the test suite. Right after performing the change, engineers are in good position to identify the obsolete tests, when compared to a spring cleaning approach done later. We believe it is more efficient, and therefore cheaper, to incentivise software engineers to keep the test suite clean and correct. GTPI is a good fit for approaching this problem from a strategic perspective, given the large number of engineers in the company and the diversity of their behaviour.

Data gathering is hard in general. Depending on the complexity of the proposed simulation model, the process engineer will need to mine from the source code repository data regarding active tests, potentially useless tests and their evolution. Also, the pay-off calculation requires data regarding how often a deleted test fails to discover a future bug and the consequences for the developer in case this happens. Obtaining this kind of data can be time-consuming or even infeasible. It might be the case that the model requires data that no one thought about collecting before the process improvement effort. The process engineer needs to start its analysis by designing a plan on how to build the dataset required for empirical game modelling, and backstop plans — like proxies — in case the organisation has not collected the required data.

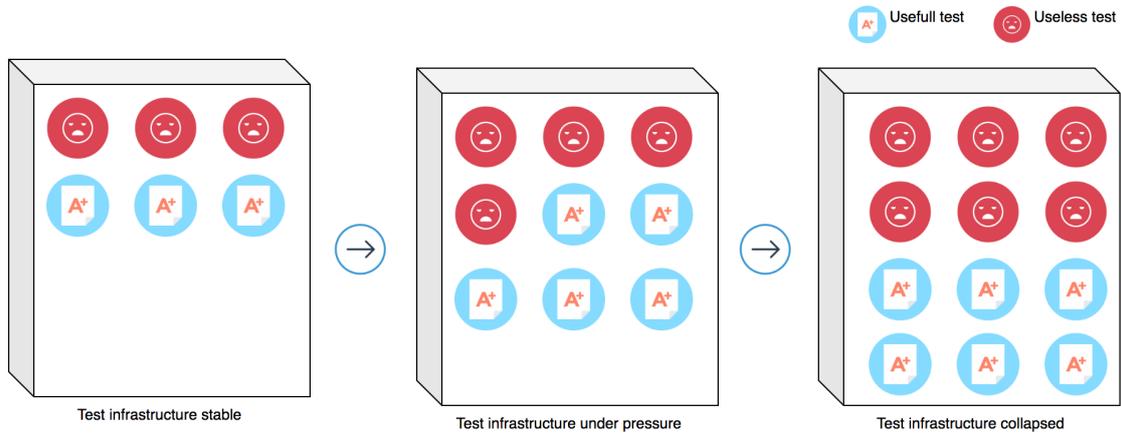


Figure 6: The tragedy of the test suite: job insecurity makes software engineers deliver more features without removing potentially useless tests. Over time, this behaviour can cause the collapse of the test infrastructure.

In the tragedy of the test suite scenario, data identifying obsolete tests is unavailable because the whole problem evaporates if the obsolete tests are known. Test prioritisation can be an effective proxy for test utility, assuming it has been empirically validated [54]. Our intuition is that the test prioritization outcome includes more useful tests than obsolete ones. Although is an imperfect proxy, it can serve as an initial baseline until more elaborate approaches are put in place.

As discussed in Section 4, after acquiring an adequate dataset the process engineer must build a simulation model of the process under study. Software process simulation is a well-known technique in the process engineering domain [9, 55], and there is a rich literature on the topic. It is critical in GTPI that the simulation model reflect the underlying process with fidelity. The empirical game improvement stage in Figure 3 requires perturbing a *validated* simulation model to obtain a process intervention that removes the undesired behaviour. Improving an imprecise model is a waste of time, so the validation of the simulation model is of the utmost importance. The validation of the model has two dimensions: technical and social. We address the *technical validation* dimension first.

In the tragedy of the test suite scenario, the simulation model needs the test execution time distribution as an input. A goodness-of-fit test can be applied to the data to see if, for example, it behaves according to an exponential distribution [48, Chapter 9]. Given that the pay-off function depends on the number of successful test executions, a statistical test can be applied to samples of the simulation output to verify if they reflect what is observed in the data [48, Chapter 10]. Simulation model validation is a hard problem, and several iterations might be needed in order to obtain a model with the required accuracy. Once the empirical game is ready, the process engineer can use any software package to calculate its NE. The obtained equilibrium also needs to be technically validated against the process data. The NE obtained from the model of the tragedy of the test suite needs to match the useless test accumulation observed in the organisation.

The empirical game model also needs to be *accepted by stakeholders*: they must agree that the model accurately captures the process. This is essential since GTPI's recommendations are learned from and justified by interventions in the model, as part of the mechanism design process. The key

challenge here is convincing the customer that the proposed model does not oversimplify.

The empirical game improvement stage requires exploring the space of potential games with a desirable NE as discussed in [Section 4](#). Building an empirical game can be computationally costly — each pay-off matrix cell needs to be simulated by several iterations — so the search space traversal needs to be done carefully. Also, a suggested process improvement intervention can fail if it is not accepted. Going back to the tragedy of the test suite, if the improved process requires a no-penalty policy for removing obsolete tests, companies can resist not punishing the unlucky developer who removed the wrong test. Cost is also an important factor: if the cost of the proposed intervention is low, convincing the stakeholders about its implementation should be easier. Given that game theory has been sparsely used in process engineering, we believe stakeholder acceptance is the biggest challenge to GTPI adoption.

The process deployment stage also proposes acceptance challenges regarding *process performers*, the people doing the actual work. The process engineer needs to ensure that the adoption of the proposed practice goes as smoothly as possible. Tool support can be crucial to this end: if the tragedy of the test suite is improved by minimising the cost of test removal, code analysis tools would help identify obsolete tests and the impact of its removal. The process deployment step needs to be monitored constantly, and perform corrective measures if needed. Without adoption, process improvement fails.

6. Conclusion

The opportunities for applying game-theoretic tools to software process improvement are boundless. We call on the community to join us in their pursuit at <https://www.researchgate.net/project/Improving-Software-Processes-via-Empirical-Game-Theory>.

Acknowledgments

This research is funded by the Dynamic Adaptive Automated Software Engineering Programme Grant (EP/J017515).

References

- [1] Talk of tech innovation is bullsh*t. shut up and get the work done – says linus torvalds, accessed: 17-09-2018. URL https://www.theregister.co.uk/2017/02/15/think_different_shut_up_and_work_harder_says_linus_torvalds/
- [2] S. McConnell, Software Estimation: Demystifying the Black Art, Developer Best Practices, Pearson Education, 2006 (2006).
- [3] F. Sarro, A. Petrozziello, M. Harman, Multi-objective software effort estimation, in: 2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE), 2016, pp. 619–630 (May 2016). doi:10.1145/2884781.2884830.
- [4] P. Butcher, Debug It!: Find, Repair, and Prevent Bugs in Your Code (nov 2009).
- [5] R. Myerson, Game theory: Analysis of Conflict, Harvard University Press, 1991 (1991).
- [6] R. Gibbons, A primer in game theory, Harvester Wheatsheaf, 1992 (1992).
- [7] I. Palacios-Huerta, Professionals play minimax, The Review of Economic Studies 70 (2) (2003) 395–415 (2003).
- [8] M. Walker, J. Wooders, Minimax play at wimbledon, American Economic Review 91 (5) (2001) 1521–1538 (2001).

- [9] J. Münch, O. Armbrust, M. Kowalczyk, M. Soto, *Software Process Definition and Management*, The Fraunhofer IESE Series on Software and Systems Engineering, Springer Berlin Heidelberg, 2012 (2012).
- [10] T. Sandholm, [The state of solving large incomplete-information games, and application to poker](#), *AI Magazine* 31 (4) (2010) 13–32 (2010).
URL <http://www.aaai.org/ojs/index.php/aimagazine/article/view/2311>
- [11] T. Sandholm, [Abstraction for solving large incomplete-information games](#), in: B. Bonet, S. Koenig (Eds.), *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, January 25-30, 2015, Austin, Texas, USA., AAAI Press, 2015, pp. 4127–4131 (2015).
URL <http://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/10039>
- [12] F. Sarro, [Predictive analytics for software testing: Keynote paper](#), in: *Proceedings of the 11th International Workshop on Search-Based Software Testing, SBST '18*, ACM, New York, NY, USA, 2018, pp. 1–1 (2018).
[doi:10.1145/3194718.3194730](https://doi.org/10.1145/3194718.3194730).
URL <http://doi.acm.org/10.1145/3194718.3194730>
- [13] M. Lavallée, P. N. Robillard, [Why good developers write bad code: An observational case study of the impacts of organizational factors on software quality](#), in: *Proceedings of the 37th International Conference on Software Engineering—Volume 1*, IEEE Press, 2015, pp. 677–687 (2015).
- [14] K. Leyton-Brown, Y. Shoham, *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*, Synthesis Lectures on Artificial Intelligence, Morgan & Claypool Publishers, 2008 (2008).
- [15] J. Nash, *Non-cooperative games*, *Annals of mathematics* (1951) 286–295 (1951).
- [16] S. Tadelis, *Game Theory: An Introduction*, Princeton University Press, 2013 (2013).
- [17] K. Holyoak, R. Morrison, *The Oxford Handbook of Thinking and Reasoning*, Oxford Library of Psychology, OUP USA, 2012 (2012).
- [18] A. Sanjab, W. Saad, T. Başar, [Prospect theory for enhanced cyber-physical security of drone delivery systems: A network interdiction game](#), in: *2017 IEEE International Conference on Communications (ICC)*, IEEE, 2017, pp. 1–6 (2017).
- [19] W. Saad, A. L. Glass, N. B. Mandayam, H. V. Poor, [Toward a consumer-centric grid: A behavioral perspective](#), *Proceedings of the IEEE* 104 (4) (2016) 865–882 (2016).
- [20] [Prison breakthrough](#), accessed: 10-09-2018.
URL <https://www.economist.com/economics-brief/2016/08/20/prison-breakthrough>
- [21] [The end of software engineering and the last methodologist](#), accessed: 22-01-2019.
URL <https://cacm.acm.org/blogs/blog-cacm/224352-the-end-of-software-engineering-and-the-last-methodologist/fulltext>
- [22] N. E. Fenton, M. Neil, [Software metrics: roadmap](#), in: *Proceedings of the Conference on the Future of Software Engineering*, ACM, 2000, pp. 357–370 (2000).
- [23] S. Parsons, M. Wooldridge, [Game theory and decision theory in multi-agent systems](#), *Autonomous Agents and Multi-Agent Systems* 5 (3) (2002) 243–254 (2002).
- [24] B. Lagesse, [A game-theoretical model for task assignment in project management](#), in: *Management of Innovation and Technology, 2006 IEEE International Conference on*, Vol. 2, IEEE, 2006, pp. 678–680 (2006).
- [25] E. Hasnain, T. Hall, M. Shepperd, [Using experimental games to understand communication and trust in agile software teams](#), in: *Cooperative and Human Aspects of Software Engineering (CHASE), 2013 6th International Workshop on*, IEEE, 2013, pp. 117–120 (2013).
- [26] G. Bavota, R. Oliveto, A. De Lucia, G. Antoniol, Y.-G. Gueheneuc, [Playing with refactoring: Identifying extract class opportunities through game theory](#), in: *2010 IEEE International Conference on Software Maintenance*, IEEE, 2010, pp. 1–5 (2010).
- [27] N. Kitagawa, H. Hata, A. Ihara, K. Kogiso, K. Matsumoto, [Code review participation: game theoretical modeling of reviewers in gerrit datasets](#), in: *Cooperative and Human Aspects of Software Engineering (CHASE), 2016 IEEE/ACM*, IEEE, 2016, pp. 64–67 (2016).
- [28] L. Feijs, [Prisoner dilemma in software testing](#), *Computer Science Reports* 1 (2001) 65–80 (2001).
- [29] N. Kukreja, W. G. Halfond, M. Tambe, [Randomizing regression tests using game theory](#), in: *Proceedings of the 28th IEEE/ACM International Conference on Automated Software Engineering*, IEEE Press, 2013, pp. 616–621 (2013).

- [30] B. W. Boehm, R. Ross, Theory-w software project management principles and examples, *IEEE Transactions on Software Engineering* 15 (7) (1989) 902–916 (1989).
- [31] H. Hata, T. Todo, S. Onoue, K. Matsumoto, Characteristics of sustainable oss projects: A theoretical and empirical study, in: *Proceedings of the Eighth International Workshop on Cooperative and Human Aspects of Software Engineering*, IEEE Press, 2015, pp. 15–21 (2015).
- [32] M. Yilmaz, R. V. O’Connor, J. Collins, Improving software development process through economic mechanism design, in: *European Conference on Software Process Improvement*, Springer, 2010, pp. 177–188 (2010).
- [33] M. Yilmaz, R. V. O’Connor, A software process engineering approach to improving software team productivity using socioeconomic mechanism design, *ACM SIGSOFT Software Engineering Notes* 36 (5) (2011) 1–5 (2011).
- [34] D. F. Bacon, E. Bokelberg, Y. Chen, I. A. Kash, D. C. Parkes, M. Rao, M. Sridharan, Software economies, in: *Proceedings of the FSE/SDP workshop on Future of software engineering research*, ACM, 2010, pp. 7–12 (2010).
- [35] M. Rao, D. C. Parkes, M. I. Seltzer, D. F. Bacon, A framework for incentivizing deep fixes, in: *Proceedings of the AAAI Workshop on Incentives and Trust in E-Communités*, 2015 (2015).
- [36] Y. Wang, D. Redmiles, Cheap talk, cooperation, and trust in global software engineering, *Empirical Software Engineering* 21 (6) (2016) 2233–2267 (2016).
- [37] O. Hazzan, Y. Dubinsky, Social perspective of software development methods: The case of the prisoner dilemma and extreme programming, in: *International Conference on Extreme Programming and Agile Processes in Software Engineering*, Springer, 2005, pp. 74–81 (2005).
- [38] J. García-Galán, P. Trinidad, A. Ruiz-Cortés, Multi-user variability configuration: A game theoretic approach, in: *Proceedings of the 28th IEEE/ACM International Conference on Automated Software Engineering*, IEEE Press, 2013, pp. 574–579 (2013).
- [39] M. Grechanik, D. E. Perry, Analyzing software development as a noncooperative game, in: *IEE Seminar Digests*, Vol. 29, IET, 2004, pp. 1–32 (2004).
- [40] N. V. Oza, Game theory perspectives on client: vendor relationships in offshore software outsourcing, in: *Proceedings of the 2006 international workshop on Economics driven software engineering research*, ACM, 2006, pp. 49–54 (2006).
- [41] V. Sazawal, N. Sudan, Modeling software evolution with game theory, in: *International Conference on Software Process*, Springer, 2009, pp. 354–365 (2009).
- [42] D. Kumar, G. Baranwal, Z. Raza, D. P. Vidyarthi, A truthful combinatorial double auction-based marketplace mechanism for cloud computing, *Journal of Systems and Software* 140 (2018) 91–108 (2018).
- [43] M. P. Wellman, [Methods for empirical game-theoretic analysis](#), in: *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference*, July 16-20, 2006, Boston, Massachusetts, USA, AAAI Press, 2006, pp. 1552–1556 (2006).
URL <http://www.aaai.org/Library/AAAI/2006/aaai06-248.php>
- [44] W. E. Walsh, R. Das, G. Tesauro, J. O. Kephart, Analyzing complex strategic interactions in multi-agent systems, in: *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, 2002, pp. 109–118 (2002).
- [45] P. Goodliffe, *Becoming a Better Programmer: A Handbook for People Who Care About Code*, O’Reilly Media, 2014 (2014).
- [46] A. Stellman, J. Greene, *Learning agile: Understanding scrum, XP, lean, and kanban*, ” O’Reilly Media, Inc.”, 2014 (2014).
- [47] A. Greasley, A comparison of system dynamics and discrete event simulation, in: *Proceedings of the 2009 Summer Computer Simulation Conference*, Society for Modeling & Simulation International, 2009, pp. 83–87 (2009).
- [48] J. Banks, J. Carson, B. Nelson, *Discrete-event System Simulation*, Prentice Hall, 2010 (2010).
- [49] S. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach*, Always learning, Pearson, 2016 (2016).
- [50] N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani, *Algorithmic Game Theory*, Cambridge University Press, 2007 (2007).
- [51] R. D. McKelvey, A. M. McLennan, T. L. Turocy, [Gambit: Software Tools for Game Theory, Version 15](#).
URL <http://www.gambit-project.org>
- [52] Q. Mi, J. Keung, An empirical analysis of reopened bugs based on open source projects, in: *Proceedings of the 20th International Conference on Evaluation and Assessment in Software Engineering*, ACM, 2016, p. 37 (2016).

- [53] C. Gavidia-Calderon, F. Sarro, M. Harman, E. T. Barr, [The assessor's dilemma: Improving bug repair via empirical game theory](#), IEEE Transactions on Software Engineering (2019). doi:10.1109/tse.2019.2944608.
URL <https://doi.org/10.1109/tse.2019.2944608>
- [54] S. Yoo, M. Harman, Regression testing minimization, selection and prioritization: a survey, Software Testing, Verification and Reliability 22 (2) (2012) 67–120 (2012).
- [55] R. Madachy, Software Process Dynamics, Wiley, 2007 (2007).