

www.ailleaders-project.eu



FindtheRightTenant



Co-funded by
the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author or authors only and do not necessarily reflect those of the European Union or the Foundation for the Development of the Education System. Neither the European Union nor the entity providing the grant can be held responsible for them.

This case study focuses on *FindtheRightTenant* – a fictional company that uses AI and credit history to decide access to housing. **It is an example of how algorithmic decision-making can lead to consumer harm and even infringe on people's rights when it does not work properly.** The case is essentially divided into two parts. First, the case study explores the history of *FindtheRightTenant* and their algorithm, explaining how it was deployed and how it went wrong, how it affected consumers and their access to housing and how it damaged the company's reputation and led to trouble for them. This part ends with the risks associated with these practices being highlighted.

The second part of the case study focuses on solutions, i.e. what a company should do to mitigate the identified risks as well as the legal norms that it ought to follow to avoid liability, as well as how it can better protect user rights and learnings that can then be applied by other companies seeking to leverage automated decision-making to decide on citizen access to opportunities in housing, credit, and other areas as well. The case study focuses on identifying the ethical risks in algorithmic decision-making related to access to housing and propose corrective measures by first identifying the ethical and legal risks associated with it so that the student can understand the issues associated with bias, discrimination, transparency, and data privacy. While the company described is entirely fictional, it is based around the experiences of an actual company that ran into similar issues. When claims are made, they are all properly backed by a source.

Goal/Purpose

The goal of this case study is to illustrate how algorithmic decision-making is integrated into business decisions related to access to housing, and to highlight the ethical and legal implications that come from the widespread use of this technology.

It does so by exploring the experiences of a fictional company, *FindtheRightTenant*, and the challenges and issues it faced when its algorithm did not work as intended and had very real consequences for people seeking to rent housing.



Expected Learning Outcomes

01

The student will be able to **identify** ethical risks in automated decision-making and scoring models, and propose corrective measures.

02

The student will **understand** the key ethical issues in AI applications for decision-making enabled by algorithms, including bias, discrimination, transparency, and data privacy.

Suggested Methodological Approach

This case works best as problem-based learning in which instructors should guide a discussion with students once they have familiarised themselves with the company and with the concepts and technology.

Topics for discussion and potential concerns are provided, but the instructors should encourage the students to think on their own and to identify other potential concerns they may have. The students should be encouraged to discuss where the problems arose from and whether the proposed solutions can be effective, as well as to come up with enhancements, improvements, or even other possible solutions.



Keywords

Algorithmic decision-making, Access to housing, transparency, information and data literacy, bias.

Introduction

FindtheRightTenant

In the Republic of Nordovia, tenant screening has become an increasingly central component of the rental application process. Rather than assessing tenants independently, **many landlords and property managers now depend on private screening companies to evaluate applicant suitability.**

These assessments often involve compiling large datasets sourced from financial institutions, credit agencies, and commercial data brokers. To streamline this process, **many screening companies have integrated AI-driven tools and automated decision-making** into their services.

FindtheRightTenant stands out as a key actor in this evolving landscape. Operating nationwide, the company provides digital tenant evaluation tools designed to help landlords make quick, data-informed leasing decisions. Since 2019, *FindtheRightTenant* has also worked in partnership with the Nordovian Social Housing Agency (NSHA), supplying tenant screening services for participants

in the National Housing Access Scheme (NHAS). This publicly funded initiative aims to expand affordable housing options for individuals and families with limited financial means by facilitating access to homes in the private rental sector.

As a result, applicants' access to public housing support increasingly depends on algorithmically-generated profiles and risk assessments produced by *FindtheRightTenant*'s systems. This growing reliance on automated evaluations raises critical concerns about accountability, fairness, and compliance with data protection laws, which safeguards individuals' rights in the processing of personal data across Nordovia and the wider European region.

Automated Tenant Screening •

Background: A General Overview of Tenant Screening

Tenant screening is an evaluation process that landlords and property managers use to determine if a prospective tenant is likely to meet their rental obligations. In other words, it is a process that allows them to identify who may meet the standards of a 'good tenant,' i.e. someone who will pay rent in time, take care of the property, and follow the tenancy contract. **The tenant screening process has become the cornerstone of modern property management, particularly on competitive markets in urban environments, where large tenants operate multiple properties.**

A typical tenant screening process includes several key components:

01

An application review, in which the tenant is asked to provide personal information, including details about employment, income, and rental history, which becomes the basis for further checks and verifications.

02

A credit history assessment, where landlords consult a credit report to evaluate the tenant's financial reliability. The landlords will look for evidence of on-time payments, debt levels, and overall creditworthiness.

03

Employment and income verification that ensures that the tenant's employment status as well as source and amount of income offer guarantees that the tenant will be able to meet rental payments. This may include checking employer references, requesting recent pay slips, or other proof of stable income such as a pension or a grant.

04

A background check, which may include reviewing eviction history, criminal records, and other available data that may help the landlords identify potential red flags. These checks can raise important questions about privacy, consent, and fairness, and, depending on the jurisdiction, they may or may not be possible.

05

Reference checks, in which former landlords or personal references are contacted so that property owners can gain insight into the applicant's past behaviour as a tenant, particularly in terms of how they took care of the property, timely payments, and the relationship with former landlords or neighbours.

Tenant screening aims to reduce financial and legal risk for landlords, but increasingly, automated means of screening have become a common practice. Automated decision-making and the use of big data also introduce concerns, raising issues such as the use of inaccurate or outdated data, which may lead to erroneous decisions; systemic bias, which may result in the system being discriminatory against certain populations; or the lack of transparency, which makes it very difficult to know how and why

decisions are made. As housing providers, policymakers, and digital service developers (the ones that create the tenant screening digital tools) engage with evolving rental markets in Nordavia and elsewhere around the world, a nuanced understanding of how tenant screening works – including comprehending its methods, benefits, and risks – becomes essential to ensure fair, transparent, and accountable housing access systems.

How Automated Tenant Screening Works

Automated tenant screening, or algorithmic tenant screening, uses artificial intelligence and machine learning to evaluate rental applications. It seeks to automate a process that is similar to 'traditional' screening described above in the sense that it also uses certain inputs to assess the suitability of a tenant. However, the use of modern digital technologies makes it possible to analyse vast amounts of data to evaluate tenants, enabling landlords and property managers to process applications more efficiently, and – assuming that the technology can work as intended – improving the chance of making the right decisions.

Automated tenant screening typically involves the following processes:

01

Data collection

Applicants fill out digital forms through which they provide personal details, information about employment, income, rental history and give consent for background checks to be performed.

02

AI assessment

Algorithms powered by AI and machine learning process this information and evaluate factors such as credit scores, rental history, employment verification, and criminal records. **AI analysis is key** and a leap from traditional assessment, as these systems can identify patterns and raise potential red flags by comparing data from a particular candidate with the insights derived from analysing large amounts of data that have been used to train the AI system.

03

Fraud detection

Advanced AI tools can also detect inconsistencies or forgeries in submitted documents, including pay stubs or identification, and intend to enhance the reliability of the screening process.

04

Decision support

The AI system generates a report or score that ascribes a **risk level** to applicants. This information can help landlords make informed leasing decisions, but, if taken at face value, and if landlords and property managers decide only on the risk score ascribed, then there is a risk that flawed systems that exhibit issues such as systemic bias lead to decisions that could be discriminatory against certain people or collectives, therefore exposing landlords to legal liabilities.

Automated screening certainly **has benefits**. It can be **more efficient**, since AI can process applications more rapidly than humans can, reducing the time between an application submission and a decision. It can be **more consistent** that human decisions, as automated systems apply the same criteria to all applications, ensuring uniformity in evaluations. **Fraud detection** can be more effective, since AI has the ability to analyse and cross-verify information and be a powerful tool for identifying fraudulent applications, as well as be more effective and faster than manual methods.

There are, on the other hand, **several challenges and considerations**, as well. The use of these systems can raise questions about the **accuracy of data**, since the effectiveness and precision of automated screening will depend on the quality and accuracy of the data analysed. **Transparency** can also be a challenge, because applicants may find it problematic to

understand how a decision was made by an opaque AI system, which, in turn, can make it difficult to contest these decisions. **Bias and fairness** can also be a significant challenge. Algorithmic decisions made with AI need to be managed carefully. Since they are trained on and use historical data to make their analysis and arrive to a decision, biases present in such data may be repeated and perpetuated, leading to **unfair outcomes** that can even constitute **discrimination**, with the potential to **systematically deny housing to certain groups in society, including those most vulnerable**.

Automated tenant screening offers significant advantages in efficiency and consistency for landlords and property managers and has the potential in aiding them to make better decisions. However, it is essential to address concerns related to data accuracy, transparency, and potential biases to ensure fair and equitable housing practices.

FindTheRightTenant Finds Trouble

FindtheRightTenant: An 'Anti'-Case-Study

An Automated Recipe for Trouble

FindtheRightTenant became essential to Nordovia's landlords and property managers, and as algorithmically-generated profiles and risk scores became central to housing decisions, FindtheRightTenant systems invited increased public scrutiny.



FindtheRightTenant's screening system relies on a wide range of data sources to generate automated reports that assign risk scores to prospective tenants. These include public records, credit information obtained from the major credit reporting agencies and the credit bureau of Nordovia, as well as other data that the company acquired from commercial data aggregators.

There are longstanding concerns related to the reliability of the datasets that *FindtheRightTenant* and similar companies have chosen to use. Many investigators have pointed to **persistent errors** in them, such as duplicate records; missing, outdated, or erroneously reported information; and issues with mistaken identity, particularly when individuals share similar names or demographic details. Such inaccuracies can compromise the integrity of the whole screening process and lead to erroneous assessments. When automated systems use flawed input data to produce what appear to be objective results, there is a risk that, if taken at face value, any outputs assessing the suitability of a candidate are, simply, wrong. **If the data used in the input is erroneous, it is more than likely that the outcome, and the decision taken as a result, is erroneous as well.**

In addition to issues related to accuracy, the data infrastructure of a system such as the one used by *FindtheRightTenant* raises concerns about the

fairness of the decisions made. Information such as criminal records of family members, employment history, a record of prior evictions, and a history of residential instability often reflect the patterns of **structural inequality** in society, which, in turn, may be rooted in Nordovia's own social problems and past issues with discriminatory housing and law enforcement practices. **Data can reflect societal issues and AI systems can reproduce and perpetuate those issues. Historical biases become embedded in the data, and thus automated systems can reproduce and legitimise these patterns under the guise of neutrality and algorithmic efficiency.**

Despite these being known issues, *FindtheRightTenant* also had a history of disclaiming responsibility for verifying the accuracy or fairness of data that it uses for its tenant screening reports. The company's practices do not include meaningful quality control on third-party data inputs it acquires, and had no risk mitigation practices as part of the way it did business. Despite being an AI-driven business, *FindtheRightTenant* had no known responsible AI practices or AI governance framework, as it claimed that its product complied with Nordovian laws related to consumer protection or personal data protection. These evident flaws caught the attention of investigative journalists in the country.



The Investigation

What Journalists Found

Several news organisations and consumer protection groups took notice of *FindtheRightTenant*'s practices. The Planet News, the National Consumer Protection League, and *The Nordovian Tech* News took the lead and highlighted several systemic problems in how tenant screening algorithms in general are designed, implemented, and regulated, but particularly singled *out FindtheRightTenant* as irresponsible.

They called out the company for not doing enough to address known issues related to **data accuracy** and the **lack of transparency**. They found that data from credit bureaus, eviction databases, and criminal justice records are often riddled with outdated, incorrect, or incomplete data. In various instances, they found out that records of arrests that never led to actual convictions and civil eviction cases that were dismissed or settled still appeared in reports, and the inclusion of these 'zombie records' resulted in **several rental rejections** that were based on misleading, false, or incorrect information.

They also found out that the company's systems were not **transparent at all** and several tenants they had interviewed had **no idea what data was used** to make a decision. They also have little to **no information about how to contest a negative screening result**. Neither *FindtheRightTenant* nor landlords routinely informed tenants when they were denied **because of an automated screening report**, with opaqueness being present in private market rentals, even if there was evidence of this problem in the Nordovian public housing system.

The journalists also investigated potential **bias and discrimination** that may be embedded in the data and algorithms present in the data that the company's systems were relying on. While they could not determine or have access to the specific databases, their research highlighted concerns about the potential **bias and discrimination** embedded in datasets and algorithms similar to

those that the company's systems were relying on, and they focused on the types of datapoints used to feed the score that feeds decision-making algorithms.

An article published by *The Planet News* explained how using criminal records, eviction history, and address-based risk scoring **disproportionately affected Nordovian minorities** – especially peoples of colour and migrants – due to **historical patterns of housing discrimination, segregation, and income inequality**. Given the fact that journalists and watchdogs pointed out that *FindtheRightTenant* had no meaningful bias mitigation and data governance procedures, they raised concerns that what looked like a 'neutral' technology that would lead to better decision-making and 'finding the right tenant' was, in fact, an automated process that carried the risk of **perpetuating and legitimising structural racism and discrimination** through those decisions.

This led to criticisms by consumer organisations about how *FindtheRightTenant* marketed its algorithmic product. They claimed that they complied with all relevant laws, but at the same time claimed that they had no responsibility for the accuracy of the data used. Both consumer organisations and the press found this contradictory and highlighted that it raised questions regarding accountability, **especially when renters were denied access to housing based on *FindtheRightTenant*'s automated reports**.



From Journalistic Investigation to Legal Action

As a result of the investigations, since 2023, *FindtheRightTenant* has been named in several lawsuits and complaints. Since it had been providing tenant screening services to both private landlords and public housing authorities, several human rights advocates and consumer organisations – and eventually also the government – sued the company.

Contracts with local governments were terminated as a result of the outcry, including a very large one with a local government agency that handled a housing rental aid programme in Nordovia's capital, making the *FindtheRightTenant*'s legal troubles very public. It soon made national and international news



The National Consumer Protection highlighted that *FindtheRightTenant* had been accused to adequately vet the accuracy of third-party data sources, and evidence has been presented in court that they did not do any vetting, even when data sources were widely known to produce frequent errors. According to the *Planet News* reporter Kent Lois, "Several of the lawsuits filed accused *FindtheRightTenant* of systematically producing reports that conflated individuals with similar names, misreported court outcomes, and presented data in ways that lacked any contest. For example, in relation to arrests, they did not include information about dismissed charges or resolutions of lawsuits."

The fact that the company got taken to court gave further tools to journalists. This led to a landmark investigation by the newspaper into the tenant screening industry in Nordovia. Facing limitations in accessing tenant data directly due to privacy laws, they turned to an innovative methodology, namely analysing federal lawsuits related to tenant screening companies.

The fact that *FindtheRightTenant* also worked with governmental agencies gave investigators an edge and, ironically, **AI tools did too**. Journalists from *The Planet News* used natural language processing and custom-built data tools to scrape and sort thousands of court records to identify patterns in consumer complaints and legal outcomes. They cross-referenced these cases with Access to Information requests to Nordovian housing authorities, which revealed the widespread use of companies like *FindtheRightTenant* in several public housing programmes

The Very Human Cost of Automated Legal Screening

How Industry Malpractice Affected the Most Vulnerable

Reporting done by The Planet News emphasised how opaque and difficult it is for renters to challenge decisions made by these systems, and how rare it is for any company to be held accountable for faulty or biased outcomes. Their investigation exposed Nordovians to a **troubling picture of the automated tenant screening industry**: a powerful, data-driven system that operates with minimal transparency, limited accountability, and serious consequences for vulnerable renters.

One case documented by the newspaper involved a woman named Luisa Sanchez who was denied housing after her tenant screening report had mistakenly included a criminal record belonging to someone else with a similar name, a Louisa Marquez. A migrant to Nordovia that had been legally in the country for almost a decade, she was repeatedly rejected from apartments and faced significant hardship as a result, despite having no criminal record herself. Stories such as this illustrate how data mix-ups and automated systems can wrongly label individuals as high-risk tenants and **severely impact their lives**.

Another man, whose name was not revealed, was flagged for a criminal record that belonged to someone else with a similar name, leading to repeated rejections. Another renter had a dismissed eviction case listed in their file, despite never having been evicted. **Both faced long delays and hardship in resolving these errors, highlighting the damaging impact of flawed automated systems.**

News articles also highlighted other renters whose lives were disrupted due to tenant screening errors. For instance, individuals were denied housing as automated reports had mistakenly listed eviction filings or criminal records that did not belong to them. These errors often stemmed from common names or outdated data. **Many victims struggled to get corrections or even understand why they were denied access.**

FindtheRightTenant, as one of the largest providers of automated tenant screening in Nordovia, served as a case study in how these systems can fail – and how such failures can harm real people seeking stable housing, **particularly when companies are reckless and do not adopt any guardrails or AI governance frameworks**. The work of journalists and advocates continues to shine a light on these practices in Nordovia today, pushing for reforms that prioritise fairness, accuracy, and transparency in housing access. Meanwhile, *FindtheRightTenant* still faces legal trouble, and while it remains in business, it is in serious financial trouble.



Guidelines for Instructors •

On the Case Study

First of all, a disclaimer: While the *FindtheRightTenant* case is based on actual journalistic investigations in the real world, this case is entirely made up for the purpose of this exercise.

First, you could start with a **discussion on AI-driven credit scoring and AI-driven financial services**, discussing the **merits of automation** and the **potential of using different types of datasets to service various populations and focus on inclusion**.

After learning about the **credit approval process**, you should **focus on the use of personal data for these kinds of decisions**. The following public sources were used to create the anti-case study, but any similarities to actual companies and living or deceased persons is purely coincidental. The instructors should make the students aware of this fact during or after they have gone through the case study and the discussion depending on how valuable the instructors may feel knowing or not knowing it beforehand is. Personally, the author of this anti-case study feels that what is more valuable is to give the students access to The Markup's actual investigation at any point (see the first provided link), as it is very good work and will help them further understand the issues at hand. When? Well, that is up to you!

You could discuss questions such as:

- What can a company do to mitigate the identified risks as well as what are the legal norms that it ought to follow to avoid liability?
- How can the company better protect user rights and learnings that can then be applied by other companies seeking to leverage automated decision-making to decide on citizen access to opportunities in housing, credit, or other areas?



The case study gives you the opportunity to identify the ethical and legal risks in algorithmic decision-making related to access to housing, as well as to propose corrective measures by first identifying the ethical and legal risks associated with it. That is what the **bold font** through the text is for. Your goal should be for your students to understand the issues associated with bias, discrimination, transparency, and data privacy.

This anti-case-study shows how a company such as this one ought **not to behave**, so part of your discussion could be about **what a company such as *FindtheRightTenant* should have done differently** and, if indeed, **whether companies that provide automated tenant screening services should even exist**, and if so, **what the best way of profiting from the advantages of technology** is while taking

This case study focuses on *FindtheRightTenant* – a fictional company that uses AI and credit history to decide access to housing. **It is an example of how algorithmic decision-making can lead to consumer harm and even infringe on people's rights when it does not work properly.**

The case is essentially divided into two parts. First, the case study explores the history of *FindtheRightTenant* and their algorithm, explaining how it was deployed and how it went wrong; how it affected consumers and their access to housing, and how it damaged the company's reputation and led to trouble for them. This part finalises by highlighting the risks associated with these practices. The second part of the case study focuses on solutions and this is where these instructions and the role of the instructors should become more active. since AI-driven models require datasets not only for training, but also for users to input personal data as demonstrated by **the simulation on page 8 of this case study**.

measures for mitigating the potential risks and tradeoffs. The instructors could encourage the students to **come up with methods and measures** to do so (see the suggested readings). Another avenue for discussion is **what those that buy a product such as the one *FindtheRightTenant* offers could do** – landlords, property managers, or public housing authorities. What is their role? Are they partially responsible or co-responsible? What can and what should they do?

Sources used:

- Kirchner, L., & Goldstein, M. (2020). How we investigated the tenant-screening industry. The Markup. <https://themarkup.org/show-your-work/2020/05/28/how-we-investigated-the-tenant-screening-industry>
- TechEquity Collaborative. (2022). Tech, bias, and housing initiative: Tenant screening. <https://techequity.us/2022/02/23/tech-bias-and-housing-initiative-tenant-screening/>
- National Consumer Law Center. (2022). New report examines how abuse and bias in tenant screening harm renters. <https://www.nclc.org/new-report-examines-how-abuse-and-bias-in-tenant-screening-harm-renters/>



Further Reading

- Asplund, J., et al. (2020). Auditing Race and Gender Discrimination in Online Housing Markets. Proceedings of the International AAAI Conference on Web and Social Media, 14(1), 123–134.
<https://ojs.aaai.org/index.php/ICWSM/article/view/7276>
- Baracas, S., Hardt, M., & Narayanan, A. (2019). Discrimination in the age of algorithms. *Journal of Legal Analysis*, 10, 113–174.
<https://doi.org/10.1093/jla/laz001>
- boyd, d., Levy, K., & Marwick, A. (2014). The networked nature of algorithmic discrimination. Open Technology Institute.
<https://www.danah.org/papers/2014/DataDiscrimination.pdf>
- Desmond, M., & Wilmers, N. (2021). Racial discrimination in housing: How landlords use algorithms and home visits to screen tenants. *American Sociological Review*, 86(5), 781–809. <https://doi.org/10.1177/00031224211029618>
- Kirchner, L., & Goldstein, M. (2020, May 28). Access denied: Faulty automated background checks freeze out renters. *The Markup*. <https://themarkup.org/locked-out/2020/05/28/access-denied-faulty-automated-background-checks-freeze-out-renters>
- Kirchner, L. (2020, October 6). When zombie data costs you a home. *The Markup*. <https://themarkup.org/locked-out/2020/10/06/zombie-criminal-records-housing-background-checks>
- Rosen, E., & Garboden, P. M. E. (2023). Algorithms and racial discrimination in the US housing market. In *The Routledge Handbook of Housing Policy and Planning* (pp. 312–324). Routledge. <https://doi.org/10.4324/9781003365877-24>
- Wang, X. et al. (2024). Algorithmic discrimination: Examining its types and regulatory measures with emphasis on US legal practices. *Frontiers in Artificial Intelligence*, 7, Article 1320277. <https://doi.org/10.3389/frai.2024.1320277>
- So,W. et al. (2022). Beyond Fairness: Reparative Algorithms to Address Historical Injustices of Housing Discrimination in the US. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 123–134.
<https://doi.org/10.1145/3531146.3533160>
- Žliobaitė, I. (2017). Measuring discrimination in algorithmic decision making. *Data Mining and Knowledge Discovery*, 31(4), 1060–1089.
<https://doi.org/10.1007/s10618-017-0506-1>

