



Microsoft's perspective on Responsible AI in Financial Services

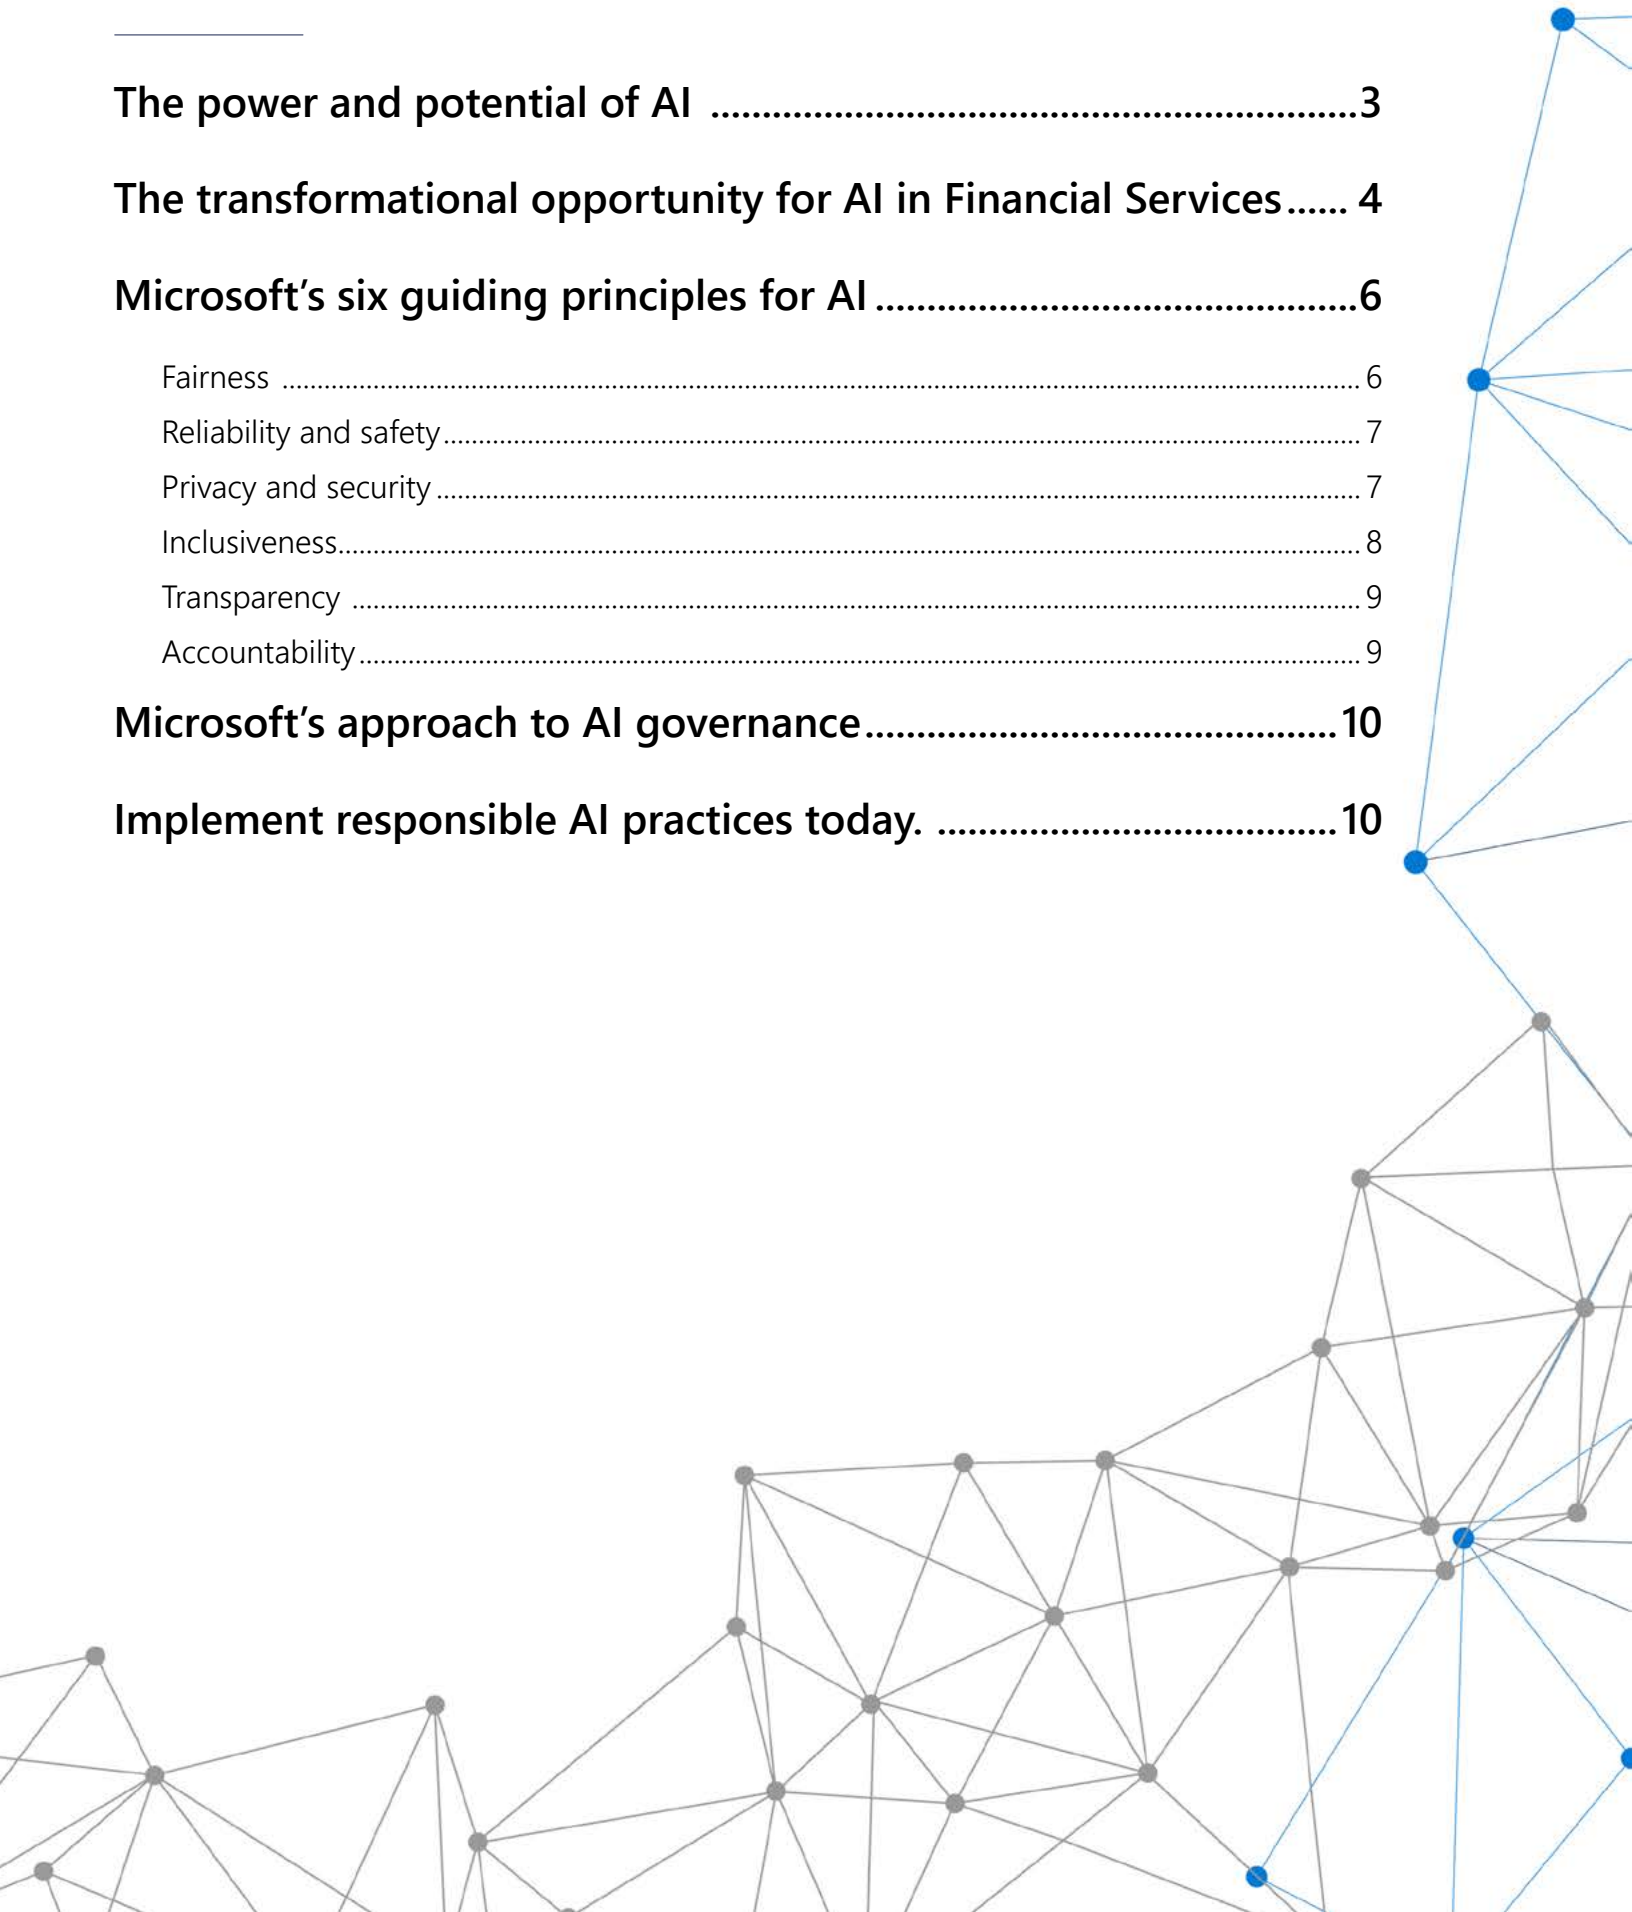
An approach to help you build trusted and respected AI

Daragh Morrissey, Director AI, Microsoft Worldwide Financial Services
Nick Lewins, Financial Services Lead, Microsoft Research



Contents

The power and potential of AI	3
The transformational opportunity for AI in Financial Services.....	4
Microsoft's six guiding principles for AI	6
Fairness	6
Reliability and safety.....	7
Privacy and security	7
Inclusiveness.....	8
Transparency	9
Accountability.....	9
Microsoft's approach to AI governance.....	10
Implement responsible AI practices today.	10





The power and potential of AI

Financial institutions are facing three major forces of change: rising customer expectations, technology advancement, and regulatory scrutiny. Customers already expect personalized, digitized experiences from other industries such as retail and media, and now expect the same from financial services. In order to compete and stay relevant, traditional financial services organizations need to understand and adopt breakthrough technologies that, by the end of the next decade, will be ubiquitous.

AI is one of the transformational technologies enabling this industry-wide shift. With the ability to ingest and analyze vast amounts of data, AI has the power to transform critical business functions, from helping financial organizations fully understand their customers—a key indicator of revenue growth and profitability¹—to more quickly and efficiently identifying fraud and security breaches. These benefits lead to many companies adopting a data-driven culture in which their leaders seek to leverage data wherever possible to increase efficiency and efficacy. Given the breadth of potential applications, it's unsurprising that the market for AI has exploded, with total spend forecasted to reach US\$46 billion by 2020, with a quarter of this coming from the financial services industry.²

While AI provides significant benefits, it's also critical to understand and manage its risks. Decisions made by financial organizations—regarding loans, interest rates, insurance coverage—are some of

This paper is part one of a two-part series.



In this paper, we explore Microsoft's ethical guiding principles for AI solutions and how they apply to financial services. In this paper, we explore how to implement governance and risk management to foster the responsible use of AI.



In part two, Responsible AI in financial services: governance and risk management, we explore how to implement governance and risk management to foster the responsible use of AI.



Total spend on AI is projected to reach **\$46B** by **2020**

the most significant decisions that can be made about a person's life. For AI to play a role in these decisions, it's essential that organizations use the technology responsibly and plan for unintended consequences. This is important for both ensuring compliance as legislation evolves and maintaining trust with consumers.

At Microsoft, we believe AI must be developed and deployed in a responsible manner. By developing and adopting clear principles that guide those building, using, and applying AI systems, we can help foster responsible AI.

In this paper, we will share the six principles we believe should guide AI development and use, and how each applies to the financial services industry. We recognize that every company will have their own AI beliefs and standards, but we want to share our perspective in the hope it helps financial services organizations develop and enhance their own unique guiding principles.

The transformational opportunity for AI in Financial Services

AI presents game-changing opportunities for transforming financial services, such as credit scoring in lending, using data analytics to forecast natural disasters in insurance, and automating the tracking of regulatory compliance changes. While there are numerous ways AI can deliver value, many AI applications ladder up to three overarching benefits: transforming the customer experience, empowering employees, and providing deeper insights.

AI can be used to improve customer experiences in many ways. Firms use AI to analyze a customer's preferences, life events, and activity to predict their needs and provide more personalized services, such as relevant product offerings or tailored investment guidance. Conversational AI helps improve customer services and engagement as well. For example, [Progressive Insurance](#) extended their

To learn more about AI use cases in financial services, see the [AI Business School](#)



marketing reach and customer support by creating a digital assistant (Flo) for Facebook Messenger, which uses natural language processing to understand customer queries and respond in the company's trademark witty style.

AI also empowers employees in the financial sector to increase productivity by surfacing insights, accelerating data analysis, and automating routine processes. With a 360-degree view of the customer, machine learning algorithms can generate guidance on the next best action for each customer, enabling financial advisors to tailor their advice and upsell new business. For instance, VeriPark (a Microsoft partner specializing in serving the unique needs of the financial services industry) developed a recommendation application called Next Best Action that helps relationship managers make quicker, smarter decisions and have more informed conversations with customers.

Finally, AI provides insights across a range of business processes. For example, banks can use AI to inform underwriting and credit decisions, which creates new revenue opportunities while more effectively managing exposure. ZestFinance created a transparent AI underwriting platform that offers AI models across consumer and commercial lending, mortgages, and auto loans. New ways of underwriting are being employed in health and life insurance practices to give coverage to customers without the need for invasive medical checks. AI also plays an important role in strengthening security by detecting fraud and money laundering, protecting customers from identity theft, and preventing cyber-attacks.

These are only a small subset of the many diverse AI solutions available to financial services organizations. As organizations evaluate which use cases to implement and how, it's crucial for them to plan ahead from day one in order to minimize risk and foster responsible design and operation.

Artificial Intelligence – Use Case – Banking



Customer Experience

Identity

B2C B2B Digital Agents

Digitized Customer Service

Targeted Offers/ Next Best Action

New Banking Products
Powered by AI



AI Powered Employees

Relationship Managers

Wealth Advisory

Risk and Compliance

Robotics Process Automation

Financial Forecasting



AI Powered Insights

Identity Protection

KYC Fraud Prevention

Payment/AML Fraud
Prevention

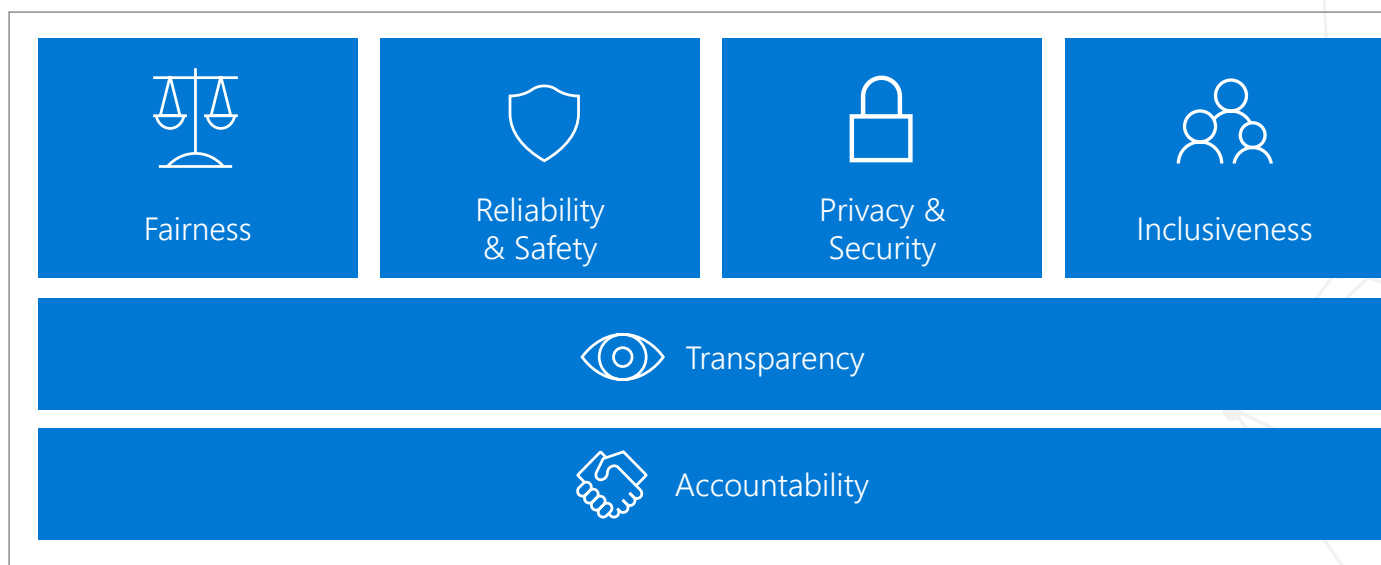
Customer Churn

Market Surveillance

Microsoft's six guiding principles for AI

Financial services organizations play a crucial role in the wellbeing of individuals, communities, and businesses. This responsibility should not be taken lightly, and financial services employees should always be cognizant of the positive and negative impact they have on these groups. As AI is being used to influence these consequential decisions, it's essential that it functions responsibly.

[We've recognized six principles](#) we believe should guide AI development and use: fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability.



Fairness

Despite all its benefits, AI could unintentionally treat people unfairly or reinforce existing societal biases. For example, imagine a large financial lending institution developed a risk scoring system for loan approvals. But because the training data reflected the fact that loan officers have historically favored male borrowers, most approved loans were for male borrowers. Without an audit, this unfairness would've persisted in the system, unintentionally adversely affecting millions.

We believe that mitigating unfairness starts with understanding the implications and limitations of AI predictions and recommendations. Ultimately, people should supplement AI decisions with sound human judgement and be held accountable for consequential decisions affecting others. It's important that developers understand how different types of unfairness can be introduced into data, machine learning models, or systems that leverage multiple models, and then use tools, methodologies, and other resources to help detect and mitigate those biases. It is also key to have robust governance processes and continually monitor models for drift or deterioration.

Reliability and safety

AI systems can make a negative impact if they do not operate reliably, safely, and consistently—even under unexpected conditions. To understand the needs for a reliable and safe solution, one needs only to imagine the role of AI in managing an organization's liquidity and underwriting functions. Such functions are systemically important to the financial stability of a bank or insurance company. Reliability of the system also plays a key role in managing financial business functions, especially in terms of managing market volatility.

As such, rigorous testing is essential during system development and deployment to prevent unexpected performance failures and ensure systems don't evolve in ways that are inconsistent with original expectations. After testing and deployment, it's equally important that organizations properly operate, maintain, and protect their AI systems over the lifespan of their use. If not maintained properly, they can become unreliable or inaccurate.

Privacy and security

As AI becomes more prevalent, protecting privacy and securing important personal and business information is becoming more critical and complex. New factors of authenticating, such as facial recognition, voice, and behavioral biometrics, are increasingly accurate and cheap, enabling new methods of payments and other activities with a lower risk of fraud.

The benefits of cost, convenience, and accuracy are clear but must be balanced with the need for privacy and security. Privacy and data security issues require especially close attention because access to data is essential for AI systems to make accurate and informed predictions and decisions. Protecting personally



identifiable information (PII) and confidential business data is paramount to an organization's success and standing with customers.

Financial services organizations must also comply with stringent privacy laws and regulations that require transparency about the collection, use, and storage of data while providing consumers with the appropriate controls to choose how their data is used.

Inclusiveness

At Microsoft, we firmly believe everyone should benefit from intelligent technology—technology that incorporates and addresses a broad range of human needs and experiences. AI can improve access to a wide range of opportunities for users with differing backgrounds, skill levels, and perspectives. And for the one billion people around the world with disabilities, AI technologies can be a game-changer, empowering them to use technologies previously unavailable.

Intelligent solutions such as real-time speech-to-text transcription, visual recognition services, and predictive text functionality are already empowering those with hearing, visual, and other impairments. These types of solutions are already gaining traction in many financial organizations. For example, many banks provide a mobile banking app where customers can conduct simple transactions through their phones. Most of these apps contain natural language processing, so users with a visual impairment can use the product by speaking instead of typing.

[Inclusive design practices](#) help system developers understand and address potential barriers in a product environment that could unintentionally exclude people. By addressing these obstacles, we create opportunities to innovate and design better experiences that benefit everyone.



3 Components of Transparency



Transparency

Underlying the preceding values are two foundational principles essential for ensuring the effectiveness of the rest: transparency and accountability. Since AI relies on extrapolated logic rather than hard-coded rules, it can sometimes function as a black box, where users don't understand how the system's outputs were derived from its inputs. But new regulations increasingly require that financial organizations provide transparency about how their AI systems work.

For example, if a bank uses an AI system to support its consumer lending decisions, it is important to examine the training data and understand how that data influences the system's recommendations.

There are three components of transparency. First, transparency relies on a foundation of traceability, with teams clearly documenting their goals, definitions, design choices, and any assumptions made in development of the system.

Second, transparency requires communication—it's important that those who are building and using AI systems are forthcoming about when, why, and how they choose to build and deploy them, as well as their systems' limitations. The third facet of transparency is intelligibility. Intelligibility means that people should be able to fully understand and monitor the technical behavior of AI systems. This understanding helps data scientists evaluate and debug models and make informed decisions about how to improve the model over time. It also helps executives, boards, employees, and customers determine how much to trust a model's predictions or recommendations.

Accountability

The people who deploy AI systems should be accountable for how those systems operate. This is especially true in industries like financial services, where AI models might help inform life-impacting decisions.

To provide oversight and guidance, we recommend establishing an internal governance system tailored to your organization's unique characteristics, culture, guiding principles, and level of engagement with AI. These governance systems can help your organization develop your own guiding principles to ensure that humans—not AI systems—are the final authority on any decision impacting people's lives.

Microsoft's approach to AI governance

At Microsoft, we established an AI and Ethics in Engineering and Research (AETHER) committee to set policies, processes, and best practices. The committee is sponsored by Harry Shum, EVP of AI and Research; and Brad Smith, President and Chief Legal Officer; and reports to Satya Nadella and the senior leadership team. It has senior leadership representation from Research, Legal, Commercial Business, and every Engineering organization. While executive buy-in is key, the committee draws input from all levels of the company to help inform their decisions.

To learn more about our guiding principles as well as the impact of AI on our future, please read our book, [The Future Computed](#).

Implement responsible AI practices today.

AI offers countless opportunities for positive change and improvements in the financial services industry, including regulatory compliance, risk, fraud detection, money laundering prevention, cyber threats, and financial inclusion. In today's connected and data-driven world, the benefits provided by AI are becoming key differentiators, leading many financial services organizations to adopt AI solutions at an increasing rate.

Nevertheless, AI technologies are nascent, evolving quickly, and present unique risks and challenges to financial institutions and their regulators. To avoid unintended consequences, mitigate risk, and minimize bias, organizations must leverage data and AI responsibly. At Microsoft, we are cognizant of these risks and are continually developing technologies and governance practices to help mitigate them. We hope our six guiding principles empower you to safely and responsibly use AI to transform your organization for the better.

In the subsequent paper, **Responsible AI in financial services: governance and risk management**, we outline potential approaches to understanding, managing, and governing AI risks in more detail. The next paper also contains key resources, tools, and methodologies you can use now to implement ethical principles of AI across your organization today.

© 2019 Microsoft. All rights reserved. This white paper is for informational purposes only. Microsoft makes no warranties, express or implied, with respect to the information presented here. This document is provided "as is." Information and views expressed in this document, including URL and other Internet website references, may change without notice. You bear the risk of using it. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

¹ Courbe, J. (2016). Financial Services Technology 2020 and Beyond: Embracing disruption. Retrieved from <https://www.pwc.com/gx/en/financial-services/assets/pdf/technology2020-and-beyond.pdf>

² IDC. (2018, H1). Worldwide Semiannual Artificial Intelligence Systems Spending Guide. Retrieved from https://www.idc.com/getdoc.jsp?containerId=IDC_P331298