

Customer journey analysis at a financial services provider using self service and data hub concepts

J.C.A.M. Buijs, R.F.M. Bergmans, and R. El Hasnaoui

Data Intelligence department, APG Algemene Pensioen Groep N.V.,
Heerlen, The Netherlands

{joos.buijs,rick.bergmans,rachied.elhasnaoui}@apg.nl

Abstract. *Situation faced:* The Dutch pension system is changing rapidly and in the future customers can choose their own pension fund. To prepare for this change APG needs to understand their customers better and further increase customer satisfaction.

Action taken: A customer journey analysis of the retirement processes using process mining analysis techniques was performed. The data contains data of 75k customers, including data regarding calls to the call center, internal processes, and website visits. The analysis was done together with business experts from different departments, including self-service training of the business.

Results achieved: Our analysis showed that 18% of customers call after completion of the retirement process. By sending a welcome e-mail once all formalities are arranged we significantly increased customer satisfaction and reduced number of calls to the call center. Another insight was that customers look up information on retirement earlier than expected, therefore we send the information package earlier, also reducing the number of calls to the call center.

Keywords: Process mining · Customer journey · Self service · Data hub · Customer satisfaction

1 Introduction

Customers expect more transparency, freedom of choice, and accountability from companies, and from the financial sector in particular. Increasing digitization and computerization has an impact on working methods and service provision. Pension funds and their customers, the participants, are facing the problems of missed indexation and possible reductions. A constant factor is the ongoing pressure to lower execution costs to minimize pension reductions for customers, while at the same time increasing customer satisfaction. Therefore the pension system is under growing pressure in the years to come, and will gradually but drastically change to afford more room for individual consumer choice.

APG (Algemene Pensioen Groep) is a financial services provider, focusing on asset management, pension administration, pension communication, and employer services. APG performs these activities on behalf of (pension) funds and

employers in a variety of sectors, but mainly education, government, and construction. APG manages € 487 billion (February 2019) in pension assets for the pension funds and customers. APG works for over 21,000 employers, providing the pension for one in five families in the Netherlands (over 4.5 million customers) [1]. APG foresees major changes in the pension scheme in the Netherlands, primarily the possibility for customers to choose their own pension fund.

Therefore APG is focusing on understanding their customers better, mainly through the ‘intense customer focus’ strategic program. Next to user group studies, customer conversations, etc. APG also used process mining [4] to analyse customer journeys, such as customer retirement.

In this paper we present the application of process mining within the customer journey of retirement, one of the key customer journeys of a pension fund. In section 2 we explain the reason for the analysis: an expected change in the Dutch pension market, combined with ongoing changing customer expectations. Section 3 presents the actions taken, for example the multi-disciplinary approach, where this paper focusses on the data analytical view of the project. We also discuss our central data store approach, which enables self-service of analytical insights within APG. The results achieved by these actions are presented in section 4. The main results directly provided for concrete actions to be taken to increase customer satisfaction. The lessons learned during the project are detailed in section 5, where section 6 concludes the paper.

2 Situation faced

One of APG’s key processes is the retirement process, a complex process taking several months from orientation, pro-active mailing, calling and all other activities to make sure the customer enjoys his well-earned retirement. For the customer this is an uncertain phase of not knowing whether retirement benefits are sufficient, which choices to make, and whether everything is arranged in a proper way. Hence customer satisfaction is key for this process.

Since APG wants to be the first choice to gain control over your financial future, it has stated ‘Intense Customer Focus’ as one of its key values, and hence customer satisfaction as one of its key performance indicators. However, one out of four calls to the call center were about the retirement process. Customers had to call themselves to ask about the status of their retirement process. These calls were complex and therefore difficult for the call agents to answer. The call center had to put the calls through to the back office. Several departments and channels were involved in this retirement process so it was not clear what the next steps in the process were, and what the customers could expect of each phase of the process. The communication from APG during the retirement process was not aligned.

The insights in subsection 4.1 show that customers called more than three times during the process with an average over half an hour call time. This means that the customer had to put a lot of effort for their retirement process. All these factors impacted the customer satisfaction in a negative way.

Additionally, the detailed insights into the success of process automation was lacking. The combination of these two factors triggered several internal initiatives to improve customer satisfaction.

One of these initiatives was the use of process mining as a technique to improve the process, because we saw that this process used several channels and systems. We realised that we needed to take the end-to-end customer journey in scope. An early version of our data hub (internally labelled ‘Data Core’) concept was available, where several systems were combined, as shown in subsection 3.4. All necessary process data from the different source systems was available in the Data Core. This data was combined with data from three contact systems, one administration system, and our customer survey system.

3 Actions taken

The overall goal of the project was to increase overall customer satisfaction, as measured through the NPS metric. In order to achieve this we needed to gain insights from the different contact points of APG with the customer over time and related to each other, in order to connect the dots between the several activities taken place in the customer journey. By correlating high or low customer satisfaction with certain customer journey characteristics we found several dials to adjust. Taking the end-to-end customer journey in scope was something we were not familiar with. We selected process mining as the tool to apply in order to get the end-to-end customer journey in view. After a successful Proof of Concept to investigate the feasibility of analysing a customer journey with Celonis [2], the process mining tool within APG, we started the analysis in 2017, scaling up to the intense customer focus program in 2018.

3.1 The objective

The main objective is to increase customer satisfaction, as measured within APG through the Net Promotor Score (NPS). As said before APG wants to be the first choice to gain control over your financial future, and therefore wants to significantly increase the NPS. This score is determined by asking our customers, on a scale from one to ten, how likely they are to recommend APG to their friends and family. People rating a six or lower are labelled as *detractors*, where scores of nine and ten are *promoters*. NPS is then calculated as $\%promoters - \%detractors$. The NPS therefore has a range of -100 (everyone is a detractor) to 100 (everyone is a promotor).

In order to improve the NPS, the idea was to find correlations between high or low customer satisfaction and customer journey characteristics. This would then result in customer journey wide suggestions, across systems and processes, for improvement regarding the customer journey retirement. Therefore it was necessary to combine business (i.e. process and customer) knowledge with data insights based on process mining of the end-to-end customer journey. In order to make sure the project would not have a one-off character we aimed at achieving

some degree of self-service in the business. Therefore we set up a program for the business where we jointly went through the following steps:

1. identify which data you need for your analysis;
2. interpret the data to derive insights and results;
3. formulate improvement actions based on these insights and results;
4. evaluate and monitor the effects of our improvement actions.

3.2 Setting up a multidisciplinary team

This initiative was the first time that the improvements were set up chain-wide instead of one department-focused. Therefore, in 2017, we set up a multidisciplinary agile team consisting of data engineers, data scientists and marketeers, the Data Insights Team. The DIT team prepared the data, tools and analysis and obtained initial insights. When the Intense Customer Focus initiative started in 2018 the DIT team joined to cover the data track within the initiative. The DIT team started to work more closely with the retirement and process experts, as well as employees from the front office, such as the call center, to obtain insights from the data analysis. One of the goals also became to teach the business experts to use and interpret the provided Celonis dashboards themselves to enable self service. This is part of the overall objective of the data intelligence department of ‘making APG data driven’.

3.3 Data collection

During the first Intense Customer Focus data session we jointly identified the customer journey on a brown-paper and subsequently identified the touch points for the customer. Using the touch points and processes we build a list of data sources to be included in the process mining analysis.

The main data challenge was to align the definitions of case and event over systems. Although it seems logical to take a customer as a case, customers can also retire twice if they retire only part-time. However, splitting the data over the two retirements was not feasible, as for instance it would become hard to assign a call to a specific retirement instance. Having the customer as the case notion, we were able to extract data from several systems and connect them together. Another issue was choosing the correct level of detail. For instance, we could include each individual process step in the customer journey, or only the executed processes without the individual step. Similarly we had each click on the myFund environment available. In the end we decided to use the executed processes, and not process steps, and myFund sessions, and not individual clicks, as events. Aspects such as the expected insights, but also the performance of Celonis, were considered.

In scope were nearly 75k customers who retired between 01-06-2015 and 01-06-2017. Per customer data is collected one year before retirement and half year after retirement:

1. Customer-behaviour through various channels like the online portal (my-Fund) (685k visits), telephony (171k calls) and email (7k emails).
2. Data about the life-event retirement (75k lifevents).
3. Personal one to one conversation (3k appointments).
4. Process logging of the pension administrative system (360k process instances).
5. Demographic data like income, gender, marital status, working sector (20+ facts for each customer).
6. Pension choices like part-time retirement (77k choices).
7. Customer satisfaction data (3k completed surveys).

In total this resulted in 75k cases with more than 20 case attributes, and over 1.3 million events in total. For APG this project was the first time where such a large amount of data was combined and analyzed at once using process mining.

3.4 Data Core

This project also again showed the usefulness of the data core, which is the ongoing data hub implementation of DI. The main idea of the data hub, as shown in Figure 1, is to provide one canonical data access layer to data in different systems within the company. The data core therefore copies data from the source systems, and translates this to objects in the presentation layer, which tools such

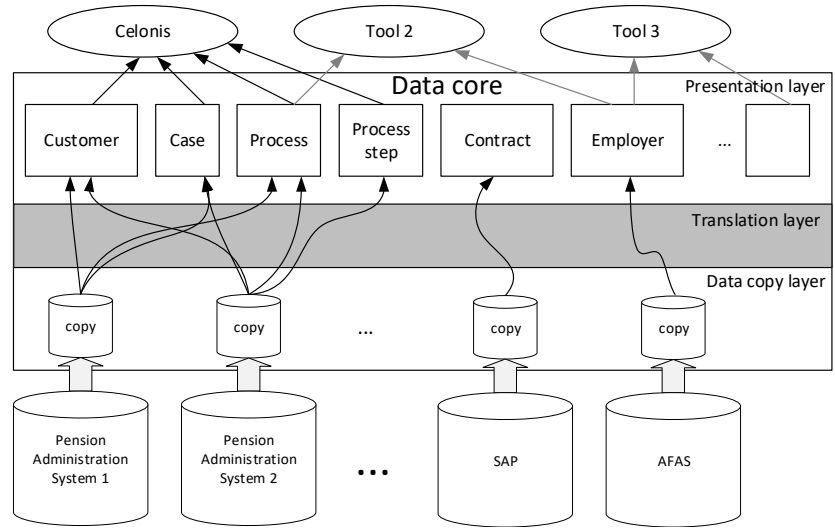


Fig. 1: The data core setup currently under construction within APG. Within the data core data is copied from source systems, translated, and presented as one canonical data source, including clear definitions of objects and attributes.

as Celonis then access. The presentation layer contains data from several systems, presented in well-documented objects (client, case, process, employer, etc.).

At the execution of this analysis phase an early version of the process part of the data core was available. This allowed us to extract all executed processes for the customers that were in scope, irrespective of the system where these processes ran. This saved us a lot of data preparation time and effort, which could now be spend analysing the rich dataset obtained.

3.5 Identifying the customer journey: exploration of facts and figures through self-service

Firstly we created a dashboard which shows the real customer journey and delivers general facts and figures. We did not do this on the traditional method where the data scientists analyse the data and build the dashboards which are then show to the business. This time we wanted, and needed, to make a co-creation through self-service. Therefore we deliver the data and then build the dashboard together with the business to deliver the insights they need. The business received practical training to understand how to interpret the Celonis dashboard and also to understand what Celonis was capable of.

We had four meetings in total where we gained several insights from the dashboards. During a meeting we worked short-cyclic and incremental. The business identified what they wanted to see on the dashboard. Most of the insights were developed ‘live’ in the dashboard during the meeting with the business watching and steering. The more difficult dashboard elements were prepared for the next meeting. This setup allowed quick hypothesis testing, leading to new hypotheses that could be tested.

The developed dashboards present facts and figures from different angles. Figure 2 shows telephone statistics such as number of calls per customer and (average) call times. This is an important aspect to consider as it related to customer satisfaction. The chart shown in Figure 3 shows the number of people retiring per month, showing clear seasonal patterns. The chart shows a clear seasonal pattern, which helps in interpretation process performance statistics correctly, knowing that the process contains more or less cases. A view on the customer-journey process flow is shown in Figure 4, which shows the complexity of the customer journey. For example, it shows how the different old and new processes are triggered, as well as when phone, myFund and mail contact occurs. Other developed dashboards showed demographic data (gender, salary range, working sector, etc.), contact channel use (comparing phone, e-mail and myFund usage), customer satisfaction, throughput-times, costs, etc.

4 Results achieved

Together with a broad delegation of the business, marketing, back-office and front-office we identified improvement points based on the obtained insights. In a final session together with the business we presented the main facts and findings.

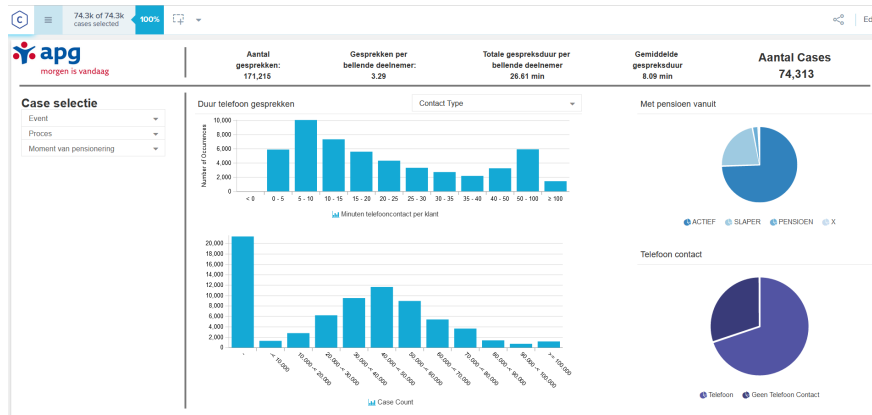


Fig. 2: Celonis dashboard showing Telephony statistics which is a key aspect of customer satisfaction.




Fig. 3: Celonis chart showing seasonality in retirement, which helps interpret process performance.

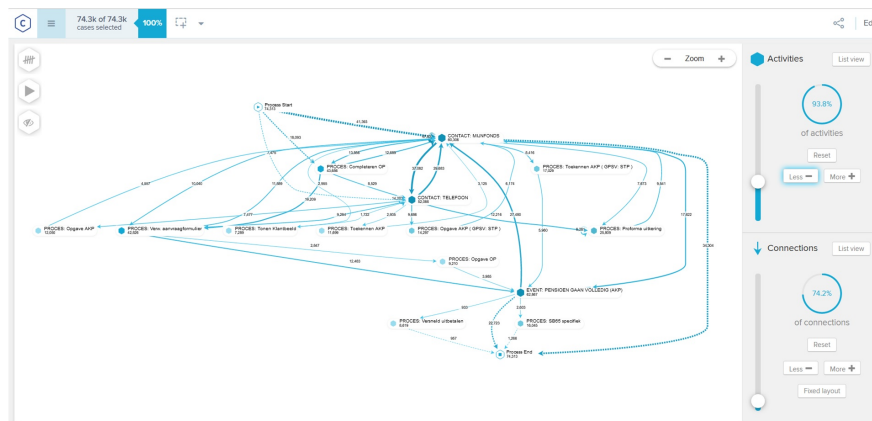


Fig. 4: Celonis dashboard showing the complexity of our customer journey flow.

With this information two key improvements were identified. The business itself took the lead into implementing these improvements.

4.1 Insights

Together with the business it was immediately clear that there was not one customer journey but many different customer journeys: more than 10k different journeys (i.e. orders of processes, letters, phone calls) were identified. One of the key findings (partly based on Figure 2) is that around 70% of the customers calls at least once for help with the call center, 21% of these customers are connected to the back-office. More than 80% of customers uses their own portal (myFund). Roughly 33% retires before the expected retirement date. On average a customer calls 3,29 times with a total call duration of 27 minutes per customer. Another observation was that 49% of visitors of the myFund environment call directly after their visit to the website. This was due to a (known) bug in the website, where the data showed the actual impact of the bug, which was far larger than expected. We also identified unclear letters send to customers resulting in a call, up to 45% for a specific letter. We also observed that 27% of customers called after receiving a requested offer from the back-office.

These insights were the starting point of useful discussions like ‘is calling three times on average good or bad?’ and ‘Why does 80% visit myFund, but far less people use the digital retirement process on myFund?’, etc.

4.2 Improvements in the retirement customer journey and communication

One of the main surprises was the observation that a lot of people call after the retirement process is completed. In general the administrative part of the retirement process is completed within three to five months before the actual retirement date. In these months there is no communication anymore with our customers. The data showed that our customers call us after the completion of the retirement process, which occurred for 18% of our customers. Further analysis revealed that most customers worry if everything is really arranged and when the first payment can be expected. Therefore we changed the process such that an email is sent 30 days before the first payment. In this message we:

1. welcome them in their retirement-life, recognizing that a new important phase in their life has started;
2. assure our customers that everything is arranged for their retirement;
3. answer practical questions that could arise, for example when the first payment can be expected.

We measured whether our customers appreciated this e-mail through A/B-testing. Almost 90% opened the email and from this group 94% gave the e-mail a ‘thumbs up’. The customer satisfaction on the end-to-end process, i.e. the transactional NPS, increased by 16 points. This e-mail was therefore considered a success and was implemented.

Another aspect we noticed was that in general our customers start to explore the possibilities of retirement 9 months before the expected retirement age. However, we send an information-package about the retirement possibilities 6 months before. Which means the information is required at an earlier point. We measured that 20% of our target-group (35k) call before they receive the first information letter, so we can avoid up to 7k calls. Therefore we changed to process to send the information-package 9 months before the expected retirement age.

4.3 Self-service introduction to the business

The self-service concept was first applied in the retirement case. We see the need that internal-clients of the DI-department want to create their own insights and don't want to be dependent of the DI-department. To answer their, sometimes simple, hypotheses we taught the business how to create their own insights by using or even adjusting Celonis process mining dashboards, hence, self service. However, this was for now too ambitious for this novice group, as we assumed a 'one size fits all' approach (see also lessons learned in subsection 5.1).

Through the sessions we did create awareness of the usefulness and power of process mining and taking a data driven approach for defining improvements. This resulted in several usages of process mining in the business, one of them is the creation of a Celonis-taskforce within the business itself, to get actionable insights and define process improvements. This contributed to the goal of DI to make APG more datadriven.

4.4 A vision on process mining within APG

The retirement case created pull from the business. The usefulness of process mining is now well known within APG and we get more and more process mining requests in order to improve processes or solve other types of problems. For example the customer journey of becoming a new customer and another journey




Fig. 5: Distribution of effort scores (x-axis) over the occurrence (y-axis) between the groups with less than two calls (group 0), and a group with two or more calls (group 1) (exact numbers are censored).

of divorcing are currently being analysed. Process mining, including self service training, is now also used for other purposes like SLA-monitoring, consistency-checks, and fraud (risk) detection.

This led us to develop a vision and roadmap towards 2025 and what operations research (OR) could bring to APG, next to our other analytics domains visualization and machine learning. The roadmap provides a clear vision both within APG, as well as within DI because data analytics experts and data scientists commit to one of three roadmaps to build their expertise. Within the OR roadmap we currently recognize four dimensions in which we want to grow:

1. From offline analysis, via prediction to recommendation of best actions to take;
2. From operational process analysis, via tactical to strategic level analysis;
3. From data on process steps via more detailed case data to very rich data;
4. From one-off analysis, to more and more continuous self-service.

The case study presented in this paper was offline, on operational processes, but on very rich data and aimed to be a continuous analysis. The ideal setting would be that all of APG, from domain experts to the board of directors make use of the same (set of) dashboards, build on the same data, and hence are aligned.

In order to achieve this we take both a top-down as well as a bottom-up approach. By organizing regular open, and non-project related community sessions, we foster the bottom-up interest. During these sessions we discuss latest features and techniques, but also let people experience process mining themselves on real data, but not directly in a project context. The main goal is increasing the self service level of the internal client. Roughly 20 people join these bi-monthly sessions in their own time, which often leads to one or more process mining analysis projects.

4.5 Follow-up data science analysis

showed that there is a relation between customer satisfaction and number and length of phone calls and/or myFund environment visits. To gain more insights a Master thesis [3] focused on combining process mining and data science methods to further analyse this observation, and make it actionable. In the Master thesis a newly developed methodology was applied on the customer journey data. It revealed for instance that customers that have to call 2 or more times provide a statistically significant lower customer satisfaction score (cf. Figure 5). It also provided a structural approach within the Celonis tool, and therefore accessible to the business users, to repeat the analysis for other findings. The master thesis therefore added concrete insights for the customer journey, as well as provide a structured approach for future analysis.

5 Lessons Learned

Based on this first large customer journey process mining project we learned two key lessons: self service is not a one-size-fits-all approach, and our data core

approach requires attention to detail. Another lesson learned is that we easily underestimate the complexity of an analysis project.

5.1 Self Service is not one-size-fits-all

What we learned with the first workshops was that self service training goes beyond providing the so called hard things as data, tools and tool-training. You have to take the soft aspects also into scope of your approach. Our colleagues have different work backgrounds, learn at different speeds, have different goals, different levels, different interests, and can be very critical of the data and technology. However, our initial approach was ‘one size fits all’, which did not fit for everyone in the group. We lost some people during the training, especially at the dashboard building part. We now use a level based maturity approach where we first determine together with the business the desired level. We then take it step by step, and at each level assess whether the level of maturity is actually met before moving on. The maturity levels focus on:

1. Correctly using and understanding the dashboard;
2. Modifying the dashboard;
3. Building a dashboard from scratch;
4. Checking the data quality;
5. Gathering and combining data.

Each level is first performed together (e.g. understand a dashboard together with data scientists), and then gradually with less and less help from the data scientists. Note that each level includes all the skill of the previous maturity levels. We are currently setting up learning goals and a curriculum to streamline the self service sessions and monitor who has been trained to a certain level.

A positive side-effect of the self service approach is that we are regularly discussing the current results with the business. We were used to analyse the data and create several cool and useful dashboards and present this to the business. Our initial analysis contained 100+ insights, of which only 4 resulted in concrete improvements. Turning insights into *actionable* insights however takes a lot of follow-up and traction. We realised too late that the self service concept also allowed us to apply focus on the hypotheses we investigate, as the business was present to indicate which line of investigation would be interesting, and which would not be. It is better to get results step by step than a big design upfront. This is something we realized too late in this project, but are applying now more often in follow-up projects.

5.2 Data core requires attention to detail

This project was also the first larger project where the early version of the data core was applied. As such we also learned with respect to certain implementation aspects. We realized that the unified process view across systems really helped. A current issue however is how to handle diversity between the systems, for instance in level of detail and existence or absence of certain attributes.

The ability to connect different objects/concepts from different systems showed the strength of the data core. However, at the same time we need to continuously pay attention that we allow these connections to be made, in other words, that objects can be linked in a concise way.

6 Conclusion

APG applied their Celonis process mining tool to analyse the full customer journey of customers retiring. During this analysis we applied our self service and data core concepts. Through the self service concept we build the dashboards jointly, teaching them to interpret and use the dashboards correctly on the way. Our data core concept allowed us to easily extract most of the data from one database where different source systems were combined and aligned.

Based on the delivered insights two concrete improvement actions were undertaken. For example, after the administrative process is completed, but before people actually retire we send them a welcome e-mail. In this e-mail we congratulate them on their retirement, but also answer some common questions, such as when the first payment can be expected. This increased NPS (customer satisfaction) by 16 points in a small test setting, which signaled a significant increase. Another improvement applied is that we now send the initial letter triggering people to arrange their pension nine months before the expected retirement date, instead of six months. This potentially saves thousands of call to our customer contact center per year.

From this large customer journey project we learned that our self service approach needed fine tuning, as it was not ‘one-size fits all’ which we kind of assumed. We also noticed the value of the data core concept, but also realized we need to be careful in its implementation. Finally we noticed a significant increase in the number of process mining project requests.

Acknowledgements: The authors would like to thank all colleagues involved in the intense customer focus project for their efforts and participation. Our thanks also go out to Celonis which provided us with a tool suitable for self service. We also appreciate the outstanding customer support and further collaboration on specific topics. All authors contributed equally to this paper.

References

1. APG. About APG. <https://www.apg.nl/en/who-is-apg>. Last Accessed: 2019-05-29.
2. Celonis. Celonis website. <https://www.celonis.com>. Last Accessed: 2019-05-29.
3. Jan-peter van der Steege. Optimizing Aggregated Performance Indicators in Service Processes using Process Mining. Master’s thesis, Eindhoven University of the Netherlands, the Netherlands, 2018.
4. Wil van der Aalst. *Process Mining: Data Science in Action*. Springer Publishing Company, Incorporated, 2nd edition, 2016.