EIGER: Auditable, executable, flexible legal regulations

Alexander Bernauer
alexander.bernauer@pwc.ch
PwC Switzerland
Zurich, Switzerland

Richard A. Eisenberg
rae@richarde.dev
Tweag
Paris, France

Abstract

Despite recent advances in communication and automation, regulations are still written in natural-language prose, subject to ambiguity, inconsistency, and incompleteness. How can we craft regulations with precision? Our solution is embodied in EIGER, a domain-specific programming language embedded in Haskell. A domain expert pairs with a software engineer to write regulations in EIGER. The domain expert needs only to *read* and *audit* the code, but not write it. A first, limited, user study suggests that this works well in practice because EIGER code mostly looks like Excel formulas with simple SQL queries.

EIGER forms the kernel of a new strategy to deliver value to clients in our professional services business with increased automation and precision. The framework is *executable*: with client data, we can use EIGER both to deduce how best to adapt to a new regulation and then maintain compliance. This paper reviews the design of EIGER and walks through its implementation. To preserve a straightforward surface syntax but with monadic semantics, we have leveraged advanced features, including *GHC.Generics*, the new OverloadedRecordDot extension, and a novel approach to performing class instance selection at run-time.

1 Regulations are *so* 20th century

In 2011, Marc Andreessen famously realized that "software is eating the world" [Andreessen 2011], transforming business, government, and private life. Yet the domain of legislation and regulation has remained largely unaffected by this revolution. Thought leaders have been explaining why the tax domain, for example, is "late" and why it is "next" [Ulbrich et al. 2019]. Triggered by the global financial crisis of 2007–2008 and fueled by "rising pressure to secure revenues" and the "challenge oactual f supervising and serving an increasingly complex and fast-paced world", authorities have responded with more and more complex laws and regulations and with "significant [investments] to develop new capabilities".

However, society has not yet reaped all expected returns: regulation drafting remains a manual, ambiguous, error-prone process, and regulators still have insufficient insight into systemic risks. Instead, regulated entities must each ensure compliance separately, duplicating the work necessary to convert regulation text into an executable specification, incurring high costs in the process [Hammond and Cowan 2021]. Other consequences include high mitigation costs for authorities, frequent lawsuits to address the deficits of the legal prose, as well as a general corrosion of the rule of law.

So, imagine instead a future when legislation in available in prose *and* in code. Critically, this code must be readable by domain experts such as lawyers and accountants. It must have precise, unambiguous, and executable semantics. The respective publishers can use well-established software engineering techniques, such as version control, automated testing, and formal verification. Checking and ensuring compliance with a new regulation can be as automated as software updates.

We are not the first to pursue this line of thinking [e.g., Fokkenrood 2011; Merigoux et al. 2021a,b; Mohun and Roberts 2020; Morris 2020; PricewaterhouseCoopers 2022]. Influenced by our background in the professional services business, however, we offer a new spin. We consider it unrealistic for individuals to become experts in both a legal domain and in software engineering. Instead, we propose a pair-programming approach: one engineer writing code and one domain expert providing requirements and verifying the implementation. Accordingly, one key criterion is whether the code is auditable by domain experts, not whether they can write functioning code. Not only do we think this approach is more realistic, but it also allows us to design an *embedded* language for regulations, making the development of such systems significantly cheaper and faster.

We have designed and implemented EIGER, a Haskell framework and domain-specific language (DSL) for encoding legal regulations. EIGER allows a heterogeneous database of records with fields whose values can depend on each other. Rules for the computation of those values can be written in any order and are written in a straightforward, pure functional style, familiar to users of Excel. When working with a dataset, EIGER accepts inputs with missing information or local overrides of fields that could otherwise be calculated; it allows users to query what extra information must be entered in order to compute a particular result. EIGER

arXiv:2209.04939v1 [cs.PL] 11 Sep 2022

also supports reading and writing JSON, enabling integrations with interactive applications and visualizations, and data pipelines.

This paper presents the following contributions:

- The user-facing design of EIGER, designed for a professional services life-cycle process (§3), and including the rule-writing interface (§4), and the end-user interface (§5).
- A limited user study (§6), where consultants not involved in the development of EIGER pair with the EIGER authors to help encode actual regulations.
- A tour of the implementation of EIGER (§7), involving heavy use of type reflection [Peyton Jones et al. 2016], generic programming [de Vries and Löh 2014; Kiss et al. 2018; Magalhães et al. 2010], GHC 9.2's new dot-notation for record access (OverloadedRecordDot) [Mitchell and Fletcher 2020], and a novel approach to instance retrieval at run-time, all built around a core inspired by Mokhov et al. [2018].
- The EIGER implementation itself.[1]

We conclude with related (§8) and future (§9) work.

2 Motivating use case: the BEPS regulation

The OECD's Base Erosion and Profit Shifting (BEPS) package addresses tax avoidance by large multinational enterprises. It "represents the first substantial renovation of the international tax rules in almost a century" [OECD 2021].

This is a sizable body of regulation with significant impact world-wide. It is currently one of the major business cases of the professional services industry within the tax domain. It is also our driving use case for the design of EIGER and the accompanying pair-programing methodology.

In the subsequent sections, we take a closer look at BEPS, explain how the industry manages to adopt it today, and highlight some of the benefits of implementation as code.

2.1 In short: it's complicated

The OECD has published model rules in a 10-chapter, 70-page document of dense legal prose. Countries are expected to implement their interpretation of these rules via their respective domestic law. Implementations will vary because each jurisdiction has different judicial interpretations, tax codes, accounting standards, and monetary policies. Jurisdictions will put BEPS into force at differing times, and thus there must be transition rules affecting multinational enterprises.

The model rules relate to existing standards such as the International Financial Accounting Standards. They also introduce new concepts such as the GloBE income, the minimum tax rate, the effective tax rate, the top-up tax, the income inclusion rule, the under-taxed payments rule, and many others. It offers dozens of elections where reporting entities

[1] https://github.com/goldfirere/eiger

can choose certain classifications that affect the rules and the tax amount owed. And it provides an interpretation of ownership structures of multinational organizations, involving concepts such as ultimate parent entities, main entities, permanent establishments, flow-through entities, and constituent entities. Finally, there are multiple exemptions, exceptions, and amendments for what does and does not count towards the tax owed, based on legal properties of entities, thresholds, prior fiscal years, etc.

2.2 How we manage today

Whenever regulators publish new pieces of draft regulations, professional services go through the following life-cycle process:

1. *Investment*: Become experts on the regulation. This involves reading, interpreting, assessing, aligning, and relating, as well as participating in the regulator's Q&A process and expert discussion groups.
2. *Knowledge management*: Capture the outcome of Phase 1 in a way that allows the business to scale up the number and size of clients.
3. *Delivery* (repeated): Use the outcome of Phase 2 and apply it on client engagements.

Executing this is very challenging in practice. BEPS has a very broad and interdisciplinary scope. This has been a trend in regulations for the last decade, as today's problems require increasingly interconnected solutions. Consequently, experts from many different tax and accounting domains need to come together to align their terminology and establish their interfaces with each other.

Furthermore, while the OECD and governments are incrementally publishing details about the BEPS model rules and their implementation in respective domestic law, regulated entities need to take action today: parts of BEPS are planned to come into force as early as 2023 in some jurisdictions. Therefore, the above process is being executed multiple times incrementally, and often in parallel.

Already without the above complications, the knowledge management step can be fragile, since it is based on more prose captured in secondary documents. This propagates the problems with ambiguity, inconsistency, and incompleteness inherent in this manual, natural-language process.

2.3 Essential versus accidental complexity

Software engineers learn how to discern essential from accidental complexity [Brooks 1986]. Thus, as software engineers, we make the following observation about the problem described in the previous section: The essential complexity of the problem lies in the semantics of the regulation, trying to capture, supervise, and influence behavior in the real world. Understanding the concepts, standards and practices involved requires training; grasping their boundaries and nuances requires experience. It is that knowledge that

domain experts offer through professional services, which is where they bring value to clients.

In contrast, the lack of formal specifications of the regulation, the lack of standardized data models and data formats, and the lack of automation in the process is accidental complexity that can and should be removed, along with the corresponding costs for implementation and operation. A specification in Eiger is unambiguous and easy to read for domain experts. Standard IDE features and related tools make the regulation navigable and discoverable. Augmented with a comprehensive test suite domain experts can also interact with the regulation through precise and correct examples. Interactive tools allow for exploration through experimentation and simulation of what-if scenarios, deepening the understanding of the regulation.

Version control can be used to track changes over time and manage adaptions for individual clients via branches. Version control also allows domain experts to share and evaluate hypothetical or factual changes through merge requests. Merge requests allow for the implementation of approval processes with controls an audit trail. Automated type checking and related static analysis yields high confidence that any changes are consistent, while the accompanying changes to the test suite convey their impact. A continuous delivery pipeline can bring factual changes to all clients in a timely fashion.

These are just a few highlights. Let's have a closer look how Eiger improves the whole professional services life-cycle.

3   The Eiger-based life-cycle

It is important to note that Eiger addresses two types of users

1. *Authors*: this is an engineer and domain expert pair that writes the Eiger code.
2. *Consultants*: this is the staff on client engagements using an existing Eiger-based system for delivery.

With that, let's revisit the life-cycle outlined in section §2.2, in the context of BEPS, but this time with Eiger in mind. We will explore its benefits and its limitations.

3.1   Investment and Knowledge Management

Unless authorities start publishing BEPS and the related domestic law as Eiger code, the majority of the investment phase stays the same. Our limited user study (§6), however, suggests that it is beneficial to start capturing insights as Eiger code early. This is even more the case when turning example cases – used to probe aspects of the regulation – into executable tests. Doing so enforces a structure that leads domain experts to a deeper and more precise understanding and guides them towards asking the right questions to eliminate ambiguities, inconsistencies, and gaps.

As the authors unveil the regulation and implement it in Eiger, they are establishing the source of truth for consultants to draw from. Together with a comprehensive test suite, this forms the bedrock of knowledge management for this regulation. The terminology used in the code becomes the lingua franca for the interdisciplinary team. It is also natural for version-controlled software to evolve iteratively while managing compatibility between versions. With this, processing the regulation and domestic laws iteratively while delivering on client engagements becomes manageable.

3.2   Delivery

A client-specific implementation of BEPS consists of three parts:

1. *Data Sourcing*: sourcing data from the client's IT systems and mapping them to input data for BEPS.
2. *Application*: applying the BEPS rules to the input data.
3. *Reporting*: rendering the results in a form proper for consumption by stakeholders.

Data sourcing is a significant challenge in practice. Solving it requires an extensive survey of the client's IT landscape, the data that are available, along with their semantics and quality. Then, proper mappings between what is available to what is required need to be developed and implemented.

Eiger cannot help with most of these problems. But it can be used to implement the client-specific mapping rules on top of the client-agnostic BEPS rules. This works well since regulatory rules are, generally speaking, mapping rules and since Eiger code is extensible (as demonstrated in Section 4). With this, the solution for the client can be managed within one holistic framework, with the exception of reporting, which is out of scope for Eiger.

Section 5 will further explore how consultants use Eiger for delivery and what benefits this brings. Before that, however, we need a basic understanding of what Eiger rules look like.

4   Writing rules with Eiger

This section demonstrates the Eiger framework by looking at parts our implementation of BEPS with Eiger.

Eiger targets regulations that determine the merits of a case by applying rules that relate to facts. Example facts include monetary values, Boolean conditions, percentages, etc., but also domain-specific classifications such as the type of a legal entity. Each fact is either provided as an input value or there is a rule that explains how to determine its value based on the values of other facts. With this, facts and rules form a directed acyclic graph.

Facts are grouped based on their context. For example, the context of a fact about income is a concrete legal entity and a concrete fiscal year. In other words, the basic structure of a regulation implemented in Eiger is an entity-relationship

model, where facts are properties and entities provide the context for facts.

Let's have a look at the BEPS prototype:

```
data Jurisdiction = MkJ {
    key :: Key Jurisdiction,
    fiscal_year :: Fact Input Int,
    additional_current_top_up_tax, adjusted_covered_taxes,
        substance_based_income_exclusion :: Fact Optional Euros,
    top_up_tax_percentage :: Fact Optional Percentage, . . . }
    deriving IsRecord
```

Every record in EIGER must have a *key* field, used in storing EIGER records in a database. The declaration above additionally says that a jurisdiction belongs to exactly one fiscal year and has a set of facts, such as *additional_current_top_up_tax*. Values for these facts may or may not be present in any given input data set, as indicated by *Optional*. Note that some facts are denominated in *Euros*, while *top_up_tax_percentage* is a *Percentage*; this separation of types provides an important additional check on the correctness of our implementation of the BEPS rules.

A jurisdiction covers a range of entities:

```
data Entity = MkE {
    key :: Key Entity,
    fiscal_year :: Fact Input Int,
    jurisdiction :: Fact Input (Key Jurisdiction),
    above_the_line_taxes, adjusted_covered_taxes,
        stock_based_compensation,
        stock_based_compensation_expense,
        stock_based_compensation_deduction,
        stock_based_compensation_expense_expired,
        stock_based_compensation_deduction_expired
            :: Fact Optional Euros,
    stock_based_compensation_election :: Fact Optional Bool,
    entity_type :: Fact Optional EntityType, . . . }
    deriving IsRecord
data EntityType = InvestmentEntity | NonSpecialEntity | . . .
```

We see that an *Entity* contains facts of a few new types, *Bool*, and *EntityType*.

Let's now have a look at a rule:

```
{-Article 3.2.2 -}
instance Rule Entity "stock_based_compensation" where
    rule self = HasRule $
        if self.stock_based_compensation_election
        then self.stock_based_compensation_expense
            − self.stock_based_compensation_deduction
            + self.stock_based_compensation_expense_expired
            − self.stock_based_compensation_deduction_expired
        else 0
```

This is a rule for the fact *stock_based_compensation* of an *Entity*. This is an example of an Excel-like formula. We use informal comments to relate code to articles of the legal prose. The identifier *self* is used to access other facts from the same entity. The rule reads such that if an election of stock-based compensation is made, certain facts are aggregated to yield the stock-based compensation. If the election is not made, then there is no stock-based compensation.

More advanced rules involve aggregating over a dynamic set of facts. Here is an example:

```
{-Article 5.1.1 -}
instance Rule Jurisdiction
    "adjusted_covered_taxes" where
        rule self = HasRule $
            getAll @Entity
                `filteredBy` filter
                `select` (.adjusted_covered_taxes)
                `andThen` sum
            where
            filter entity =
                (entity.jurisdiction == self.key) `and`
                (entity.fiscal_year == self.fiscal_year) `and`
                {-Article 5.1.3 -}
                (entity.entity_type /= use InvestmentEntity)
```

This rule, resembling an SQL formula, is for the *adjusted_covered_taxes* fact of a *Jurisdiction*. The rule says that the adjusted covered taxes of a jurisdiction is the sum of the adjusted covered taxes of all entities that meet the filter criteria: the entity must be in that jurisdiction, be associated with the same fiscal year, and not be an *InvestmentEntity*.

Finally, some facts have no rules at all:

```
{-Article 5.4.1 -}
instance Rule Jurisdiction
    "additional_current_top_up_tax" where
    rule self = NoRule
```

Facts like *additional_current_top_up_tax* are listed as *Optional* in order to allow working with a dataset missing that datum, but we still have no way of computing it if it is missing.

If we wanted, we could later decide to introduce a rule for additional current top-up tax that determines the value from raw data. The rest of the regulation and its implementation would remain unaffected by this extension. Referring back to Section 3.2, this is the interface between the client-agnostic BEPS core and the client-specific mapping.

Now that we understand what EIGER code looks like, we can further explore its use and benefits.

5 Running your regulations

Given the implementation of BEPS with EIGER described in the previous section, what does this prove useful during delivery?

Typically, regulated entities adopt regulations in two distinct phases.

1. *Exploratory:* This is about understanding the opportunities, benefits, liabilities, and risks that come with the regulation. There is a strong connection of such an analysis with the entity's business strategy and its business model. The general question is: how the entity should best position itself for the time when the regulation will be in force?

2. *Operational:* This stage is about managing compliance with the regulation on an ongoing basis: controls, reporting, and investigations. At this point it is all about reliability and costs.

The following sections explore how Eiger brings benefits to users in both phases.

5.1 Exploratory

In the context of BEPS, typical questions in the exploratory phase include:

- How much additional tax are we projected to owe in each of the next five years?
- What are the options available to reduce our tax footprint?
- How much would implementing such a change cost and how much would it save?

Exploratory analysis might suggest a wide range of changes to an organization's structure. For example, it might be recommended to relocate certain entities, acquire, sell, split, or merge other entities, move staff between entities, as well as change the flow of goods and services through multinational subsidiaries. For each of these ideas, the consultants need to collect relevant input data, such as the entity's financial statements, and apply the rules of the regulation. Typically a different subset of the rules comes into play each time, so the work tends to not be repeatable. The potential savings are often in the millions, and mistakes can lead to unpleasant surprises such as heavy fines when interacting with the authorities later.

Today, such work is mostly done manually with light support from Excel sheets and related extract-transform-load and process automation tools. None of these tools provide the flexibility required to exhaust the available options efficiently and with precision.

An interactive system loaded with a BEPS implementation in Eiger, in contrast, offers a rich user journey. Users can browse the data model and the fact graph. They can also work interactively by providing a set of input data and have the system return a superset which includes values for all facts that can be computed from the provided input.

With this, users can run:

- *Simple queries*: The user can ask a regulation-as-code system about the value of any of its facts. As analysts learn more, they can tailor their queries accordingly.

- *What-if scenarios*: Sometimes the user does not yet have the data to determine if a multinational enterprise is eligible, but they still want to know what eligibility would imply. In such a situation, the user can provide an input value of *True* for the fact about eligibility and query the value for the top-up taxes owed.

- *Optimizations*: BEPS allows for several elections on how legal entities wish to be classified. Those elections have an impact on the tax owed, so it is natural to ask which election choices are most preferable. The user can apply all combinations of elections in turn, and query the tax owed in each case.

Based on our experience, none of the tools used in such client engagements today provides this degree of flexibility with high precision and short lead times.

5.2 Operational

Once the regulation is in force, certain obligations, such as reporting, have to be met. Today, this involves extensive IT projects involving consultants, business analysts, developers, testers, and operators. While this process also yields a machine-executable implementation of the regulation, the resulting code is not auditable by domain experts. This leaves a gap with plenty of opportunity for misunderstandings between business analysts, developers, and testers. The implementation is also custom-made, so its costs cannot be mutualized across multiple entities, and the interpretation of the regulation is likely to be unique. This leads to poor data quality as observed by the authorities across the regulated industry.

Most of these problems are addressed by Eiger. A single implementation can serve as the core for the production implementations of all regulated entities. It is the same well-tested implementation that has been used during the exploratory phase. For each regulated entity, there is a one-time effort to integrate the Eiger core in the respective IT landscape. It is then fed with raw data and, in contrast to the exploratory phase, regularly executes the same queries using different data once for each reporting period.

6 User Study

To validate the design of Eiger and the approach of pairing domain experts with software engineers, we have run a limited user study. We paired with a domain expert from each of three different domains and implemented a non-trivial subset of a regulation. Afterwards, we asked the domain experts the following questions:

1. How easy was it to get comfortable with the pair-programming approach?
2. What's the one thing we could add to make your experience better?
3. How high is your confidence that what you are seeing is a faithful implementation of the legal text?

4. How confident are you that you could read other people's implementation of legal text based on this approach?

5. If, in future, you were asked to advise on regulatory matters, would you use the code or the legal text as a reference? If not the code, what is missing for you to use this instead?

6. Given the prospective benefits of rules as code as an idea, do you consider the pair programming approach a worthwhile investment?

7. Based on your understanding of the code, are there any specific areas where you think the code would struggle to replicate existing regulatory provisions?

8. What are your main feedback points about the process (both good and bad)?

Section 2 above provides an introduction to BEPS. This is the first case of our user study. As part of this study, and in preparation for a client engagement, we have implemented most of chapters 3, 4, and 5. The other subjects of study are the EU's Carbon Border Adjustment Mechanism (CBAM) and the UK's Value Added Tax Regime (VAT). CBAM requires businesses to offset financially the carbon emissions embedded in the products they import from non-EU countries, while VAT is applicable to the sales of goods and services within the UK. Both choices of subject pushed the design of Eiger forward, requiring different features from those used in the BEPS domain.

We implemented a prototype for each case by pairing one of the authors of this paper as an engineer with one of three domain experts. While the verbatim feedback is available in the appendix, we present a summary here. All three domain experts report feeling comfortable with the pair-programming approach and expressed confidence that the respective implementation is an accurate representation of the respective regulation. The experts also responded confidently about their ability to read other people's Eiger code after a short introduction, eventually using the code as a reference moving forward. Some experts noted that being forced through the structured process of codifying the regulation deepened their own understanding. They also expressed surprise how quickly the pair managed to produce functional code and how mutually beneficial it is to work as a pair. Overall, the experts deemed the investment in this new approach worthwhile, and they considered the approach to be scalable.

From the perspective of the engineer, pair-programming feels very productive. We can guide the expert with directed questions to help us understand the overall structure first, and then incrementally dive into the details. As we stumble upon open questions during the implementation, we get immediate answers, keeping the train of thought going. Sometimes, we need to ask for patience when refactoring or fixing compiler errors. And, sometimes, we need to take notes to

work independently later, when some deeper thinking, research, or a larger cleanup is needed. So far, it has always been easy to loop the expert in again afterwards. Finally, it takes some additional effort to keep experimenting with code formatting and implementation techniques to optimize for readability. To do that well, engineers need a solid understanding of Eiger, which is provided in the next section.

7 A Tour of Eiger

There are two key design challenges in implementing Eiger: designing a mechanism for writing computation rules for fields in terms of other fields, and making the syntax for writing these rules straightforward enough for a regulations expert to understand. Because we want fields of Eiger records to be typed (preventing, say, the addition of a monetary amount with a percentage), we must build out some infrastructure to deal with a heterogeneous database (that is, storing many different types of Eiger record) with a variety of field types. Then, despite the complexity under the hood, we must produce an API that is easy for domain experts to read and understand.

7.1 Tracking dependencies

An Eiger rule is specified by giving a *ComputationRule t*, which is a description of how to compute a value of type *t*. *ComputationRule* has a simple definition:

```
newtype ComputationRule t =
    MkRule { runComputationRule :: EigerM t }
```

A description of how to accomplish a (effectful) task is just a monadic action. We use our *EigerM* monad, whose details are in Section 7.1.3. Despite this simplicity, we keep *ComputationRule* as an abstraction layer for our users, allowing us to change design later. For example, we might imagine constructing syntax trees describing the operations needed for the rule for further analysis. Furthermore, *ComputationRule* has a number of instances written for it to make it work in our DSL, but we do not want these instances affecting our internal *EigerM*.

Within the computations stored in a *ComputationRule*, access to data is through two functions: *accessField* and *accessAllKeysOfType*. Both of these track dependencies, as described below.

7.1.1 accessField. Below is the *accessField* function:

```
accessField :: ∀ (field_name :: Symbol) record.
    (HasFact field_name record, KnownSymbol field_name) ⇒
    Key record → EigerM (FactType record field_name)
accessField key = do
    recordFieldDependency @field_name key
    record ← getRecordFromKey key
    getFact @field_name record
```

The *accessField* function is polymorphic in a *field_name* (recall that *Symbol*s are just compile-time strings) and a *record* type; these are related by the *HasFact* class. A constraint *HasFact name record* states that the record type *record* has a field named *name*. The field will be a *Fact* whose payload is *FactType record name*, the type returned by *accessField*. *FactType* is an indexed type family [Chakravarty et al. 2005a; Eisenberg et al. 2014], which can be considered a compile-time function, computing on types. In this case, *FactType* looks at the type definition (via the *GHC.Generics* interface [Magalhães et al. 2010]) of the record and extracts out the type of the fact of the given name. The function is also constrained by a *KnownSymbol* constraint; this says only that we know the field's name at run-time. Its name is needed only for dependency tracking.

The function takes an argument of type *Key record*, which specifies which record to retrieve from the database. We will see more about how *Key*s in Section 7.2.1 and how *getFact* retrieves the data in Section 7.2.4.

Crucially, the *accessField* function also notes internally within the *EigerM* monad that a *ComputationRule* required a certain field, via *recordFieldDependency*. This is necessary to support a user mode where an analyst can query what dependencies a given field has.

7.1.2 *accessAllKeysOfType*.
The *accessAllKeysOfType* function is defined as follows:

accessAllKeysOfType :: ∀ *record*. *Typeable record* ⇒
    EigerM [*Key record*]
accessAllKeysOfType = **do**
    recordTypeDependency @*record*
    getAllKeysOfType ⟨\$⟩ *getDatabase*

This function uses a run-time type representation (retrievable from the *Typeable record* constraint) for a record type to retrieve a list of all keys to records of that type. Once again, the *accessAllKeysOfType* function remembers this call, via invoking *recordTypeDependency*. This way, an analyst can learn that, for example, adding a new *Entity* with a certain *Jurisdiction* might affect the calculation of *Jurisdiction* fields.

7.1.3 The *EigerM* monad.
The *EigerM* monad is used throughout Eiger to structure our computations. It has a mix of standard monad features. *EigerM* is:

- *State*-like over our *Database* (§7.2.1);
- *Writer*-like over a *Dependencies* monoid tracking dependencies; and
- *Except*-like over a list of errors, including missing dependencies.

One detail in this list may be surprising, in that monadic actions can result in a *list* of errors. This is unusual: a monad able to throw exceptions often throws only one!

We can explain this twist with an example. Suppose we wish to compute a total income from an unearned income and an earned income, but we actually have neither fact in our data set. We might imagine a computation that looks like this (in a more Haskell-like syntax than our DSL as demonstrated in Section 4):

getTotal compute_unearned compute_earned = **do**
    unearned ← *compute_unearned*
    earned   ← *compute_earned*
    return (*unearned* + *earned*)

However, if there is no way to *compute_unearned*, the first line here will throw an exception. No other possibility exists because there is no valid answer to return from the computation. Yet, from a user standpoint, it would be a shame not to learn that *compute_earned* is just as impossible as *compute_unearned*.

The *EigerM* monad thus has a carefully written *Applicative* instance. In its definition for $(\langle * \rangle) :: Applicative\ f \Rightarrow f\ (a \rightarrow b) \rightarrow f\ a \rightarrow f\ b$, we run the second action *even if the first one fails*. This behavior makes our instances not quite lawful: the laws for *Applicative*s state that *Applicative*s that are also *Monad*s must obey the law saying that $m_1\ \langle * \rangle\ m_2$ is identical to $m_1 \ggg \lambda x1 \rightarrow m_2 \ggg \lambda x2 \rightarrow return\ (x1\ x2)$. Yet this is not the case for *EigerM*: $m_1\ \langle * \rangle\ m_2$ performs both actions, while the right-hand side of the law performs only m_1 if m_1 fails.

It is tempting to try to bring the *Applicative* instance and the *Monad* instance back into alignment, but this is impossible. In general, computations may need *Monad*-like behavior: we might have one fact in our database that tells us where to look next. For example, our database might give us the keys of records describing the dependent children of a taxpayer; only after retrieving those keys can we access the incomes of the dependent children. Accordingly, in these scenarios, our dependency tracking is on a "best effort" basis: we can look for dependencies only with the information we have, and it is possible that providing more data for a computation will reveal more dependencies. In order to allow the possibility of conditional computations, this situation is unavoidable.

Is it problematic to have these unlawful instances? In our opinion, no. The *Applicative* combinators will be more informative than their *Monad* counterparts, but both results are guaranteed to be subsets of the true set of dependencies. Furthermore, it seems unhelpful to users to allow strict compliance with laws to stop us from reporting as many dependencies as possible.

Our situation is rather like that faced by the *Haxl* monad described by Marlow et al. [2014, Section 5.4], where their *Applicative* instance leads to better performance than their

Monad instance. Just like them, we think it is acceptable to bend the laws here.[2]

7.1.4 User interface.
The design described here supports the following client-facing functions:

get :: ∀ (*field_name* :: *Symbol*) *record*.
  HasFact field_name record ⇒
  Key record → *EigerM* (*FactType record field_name*)
getMissingDependencies :: ∀ (*field_name* :: *Symbol*) *record*.
  HasFact field_name record ⇒ *Key record* →
  EigerM ([*String*], [*TyCon*])

The *get* function retrieves the value of the field named *field_name* from the record keyed by the key provided. It uses rules as necessary to compute any field that is not present in the database and updates the database as it works.

The *getMissingDependencies* function returns lists of all facts that are needed to compute a target fact (but are not present), as well as a list of all record types that were queried via *accessAllKeysOfType*. The function is useful against a data set that is missing some input facts. For example, if we want to know the total income of an individual and know their earned income but not their unearned income, this function would say that we must input the unearned income. The list of record types consulted by *accessAllKeysOfType* is important because an addition (or deletion) of such a record might change the result, as well.

In addition to *getMissingDependencies*, we can imagine further queries of this form, including listing all dependencies (missing or not).

7.2 A typed interface
This section explores the challenges of working in a richly typed environment, and how EIGER overcomes these challenges.

7.2.1 A heterogeneous *Database*.
An EIGER data set is *heterogeneous*: it contains records of many different types, such as the *Jurisdiction* and *Entity* types of Section 4. We must store these different types in one database. Furthermore, this database must be designed without knowledge of what types it will be storing, and yet we wish to avoid extraneous run-time type checking.

The solution is to use a *dependent map* [Cook and Obsidian Systems 2020]. The dependent-map Haskell package exports the following type and functions:

type *DMap* :: ∀ *ki*. (*ki* → *Type*) → (*ki* → *Type*) → *Type*
data *DMap k f* = . . .

---

[2]The *Haxl* monad was the motivating example for the `ApplicativeDo` extension [Marlow et al. 2016], which desugars **do**-notation to use *Applicative* combinators where possible. However, our user-facing syntax avoids **do**-notation, so we do not use `ApplicativeDo`, though it would indeed work well for EIGER computations written with **do**-notation.

data *Key record* = *UnsafeMkKey* (*TypeRep record*) *UserKey*
data *UserKey* = *ExternalUserKey Text*
          | *InternalUserKey InternalKey*
newtype *InternalKey* = *MkIK Int* -- kept abstract
data *Database* =
  MkDb {*database* :: *DMap Key DatabaseValue*, . . . }
databaseLookup :: *Database* → *Key record* → *Maybe record*
getAllKeysOfType :: *Typeable rec* ⇒ *Database* → [*Key rec*]

Figure 1. Excerpt of *Database* interface

lookup :: *GCompare k* ⇒ *k v* → *DMap k f* → *Maybe* (*f v*)
insert :: *GCompare k* ⇒ *k v* → *f v* → *DMap k f* → *DMap k f*

A *DMap k f* is a finite map connecting keys of type *k v* to values of type *f v*, for any *v*. Accordingly, *lookup* and *insert* work over keys *k v* and values *f v*; these functions are polymorphic over the choice of *v*. Using *DMap*, we can map a key of type *Key record* to a value of type *record*, using the following wrapper, defined as part of EIGER:

type *DatabaseValue* :: *Type* → *Type*
data *DatabaseValue rec* = *MkDbV* (*TypeRep rec*) *rec*

We see here that a *DatabaseValue* stores both a run-time type representation of a record (this is *TypeRep record*) and a value of type *record*. The *TypeRep record* is necessary to implement *getAllKeysOfType*, as presented in the *Database* interface in Figure 1.

We see in Figure 1 also the function *databaseLookup*, which finds the record associated with a key, the critical operation in our heterogeneous database; the implementation of this function is a straightforward use of *lookup* from the dependent-map package.

Figure 1 also presents the *Key* type. A *Key* stores a run-time type representation of a record, as well as a unique identifier for the record to look up. EIGER supports both user-specified identifiers, stored as *Text*, as well as internally generated keys, represented simply as *Int*s. The elided fields of a *Database* include some mundane structures that guarantee uniqueness of *UserKey*s in the database.

A *Key* needs a type representation in order to safely implement the *GCompare* class, exported in the some package [Cook and Grenrus 2021]:

type *GEq* :: ∀ *k*. (*k* → *Type*) → *Constraint*
class *GEq f* **where** *geq* :: *f a* → *f b* → *Maybe* (*a* :∼: *b*)
class *GEq f* ⇒ *GCompare f* **where**
  gcompare :: *f a* → *f b* → *GOrdering a b*
data *GOrdering a b* **where**
  GLT, *GGT* :: *GOrdering a b*
  GEQ   :: *GOrdering t t*

Note that the result type of *gcompare* is a *GOrdering*, whose *GEQ* constructor requires that the indices of the types being compared (that is, *a* and *b*) are the same. If a *Key* stored only the unique identifier *UserKey*, there would be no way to prove to the type checker that when two *UserKey*s are equal, the types of the records they index are also equal. The *TypeRep* stored in a *Key* allows us to make this check and establish type safety.

7.2.2 Field access. Suppose we have a record *r* of type *record*. We will need a way of extracting the value of a field *name*. Recent versions of GHC provide its *HasField* class [Gundry 2017], which would seem to do the trick:

class *HasField x r a* | *x r* → *a* **where**
  getField :: *r* → *a*

In fact, *getField* @*name r* works. However, this approach is insufficient for our needs. Not only do we need to access fields, we need to *update* them after performing a computation on a field derived from others via user-written rules. GHC has no facility analogous to *getField* for updating fields, due to debates about what the best design is.[3]

The generic-lens package [Kiss et al. 2018] offers a way forward. The package exports these functions:

getField :: ∀ *f a s. HasField' f s a* ⇒ *s* → *a*
setField :: ∀ *f a s. HasField' f s a* ⇒ *a* → *s* → *s*

The *HasField'* class is part of generic-lens; it is distinct from the *HasField* class above that is part of GHC. *HasField'* instances are built using the generic representations from *GHC.Generics* [Magalhães et al. 2010].

Because we wish to use functions named *getField* and *setField* as part of the interface presented to authors writing rules (via the OverloadedRecordDot and RebindableSyntax extensions; see Section 7.3.3), we rename and repackage these functions for internal use, where *GL* is a module alias to the appropriate module from generic-lens:

  -- this includes *Generic*, required for a *HasField'* instance
type *HasRecordField field_name rec field_ty* =
  (*HasField' field_name rec field_ty*, *Generic rec*)

getRecordField :: ∀ *name rec ty*.
  HasRecordField name rec ty ⇒ *rec* → *ty*
getRecordField = *GL.getField* @*name*

setRecordField :: ∀ *name rec ty*.
  HasRecordField name rec ty ⇒ *rec* → *ty* → *rec*
setRecordField r x = *GL.setField* @*name x r*

---

[3]There is an accepted proposal [Mitchell 2019], but it has not been implemented and personal communication suggests a revision of the plan is in the works.

7.2.3 Facts. As we saw in Section 4, an Eiger record contains fields declaring *Fact*s. This type and related definitions are here:

data *FactSort* = *Input* | *Optional*

type *Fact* :: *FactSort* → *Type* → *Type*
data family *Fact sort t*

newtype instance *Fact Input t* = *InputFact* { *unInputFact* :: *t* }
data instance *Fact Optional t* = *MustDerive*
               | *AlreadyKnown t*

type *GetFieldFactSort* :: *Type* → *Symbol* → *FactSort*
type family *GetFieldFactSort record name*

We first define a *FactSort* as either *Input* or *Optional*. An *Input* fact must be specified in order for a record to be well formed; an *Optional* fact might be missing. Some *Optional* facts have computation rules associated with them, but others do not – a fact missing in the input and with no computation rule is simply unknown.

The *Fact* type is a *data family* [Chakravarty et al. 2005b]. Data families are like the more common type families in that they allow pattern-matching on their arguments, but a data instance is required to produce a fresh datatype or a newtype. Accordingly, data family types can appear unsaturated in Haskell (type families must be applied to all of their arguments), and it is possible to write type (or data) families that pattern-match on data families.

This ability to match is essential in the implementation of *GetFieldFactSort*. This type family, needed to define *HasFact* instances (in the next section), takes a record type and a field name, returning the *FactSort* of the fact stored in that field. It does this in two steps: first, *GetFieldFactSort* must retrieve the type of the field; then, it must extract out the *FactSort* from something like *Fact Input EuroAmount*. The second step requires that *Fact* be a *data* family, not a *type* family.

To implement the first step, we must have a way of accessing a field's type given its enclosing record type and the field's name. Sadly, the *HasField* interface of Gundry [2017] works via a functional dependency and is thus inadequate for obtaining a field's type from its name and enclosing record type. Instead, we must build off the *GHC.Generics* library (details of the type families are elided due to space) to finish the implementation of *GetFieldFactSort*.

7.2.4 HasFact. The fact retrieval mechanism centers on the *HasFact* constraint. We cannot, however, define this as a class directly. This is because we must write separate instances for *Input* facts and for *Optional* facts; yet whether a fact is an *Input* or is *Optional* is not apparent from the record type and field name. We thus create the internal *HasFact_* class and define *HasFact* in terms of *HasFact_*:

type *HasFact_* :: *Symbol* → *Type* → *FactSort* → *Constraint*
class (. . .) ⇒ *HasFact_ field_name record sort* **where**
  getFact :: *record* → *EigerM* (*FactType record field_name*)

```
type HasFact field_name rec =
  HasFact_ field_name rec (GetFieldFactSort rec field_name)

instance {−# overlapping #−}
  ( ..., FactType record "key" ∼ Key record) ⇒
    HasFact_ "key" record Input where ...

instance ( ..., GetFieldFactSort record name ∼ Input) ⇒
    HasFact_ name record Input where ...

instance ( ..., GetFieldFactSort record name ∼ Optional) ⇒
  HasFact_ name record Optional where
  getFact record = case getRecordField @name record of
    AlreadyKnown result → return result
    MustDerive → case rule @record @name key of
      NoRule → missingRule @name record
      HasRule r → do
        result ← runComputationRule r
        updateDatabase key $ λrec →
          setRecordField @name rec (AlreadyKnown result)
        return result
    where
      key = getRecordField @"key" record
```

The *HasFact_* class is actually quite simple: it allows retrieval of a fact from a record in the *EigerM* monad. The first two instances declared are also pretty easy. We have a special case for the "key" field, which is declared outside the *Fact* structure, as it is required for every Eiger record type. *Input* fields just return their contents.

For *Optional* fields, we first check whether the field value is *AlreadyKnown*; if so, return it. Otherwise, we look up the rule. The constraints to this instance insist that there is an **instance** *Rule record name* available, but it is possible that this instance specifies *NoRule*. If there is *NoRule*, we throw an error in the *EigerM* monad using *missingRule*. Otherwise, run the rule, updating the database with the result (to avoid recomputation).

7.2.5 JSON serialization via run-time instance lookup.
We must provide users a way of inputting a large quantity of data to seed the database, and of outputting an updated database after performing computations. EIGER thus supports serialization and deserialization with JSON. The design for such a mechanism is challenging, especially for deserialization: the library does not know the set of types that need to be parsed, and the input will be heterogeneous. Because an input may have an arbitrary number of different records of different types, each one must declare its own type. Here is an example of what an input JSON file might look like:

```
{"Corp":{"type":"Entity"
        ,"jurisdiction":"Switzerland"
        ,"fiscal_year":2022
        ,"stock_based_compensation":12345.00}
,"Switzerland":{"type":"Jurisdiction"
```

```
        ,"fiscal_year":2022
        ,"top_up_tax_percentage":0.03}}
```

The top-level structure is a JSON object, where the keys of the object will turn into keys in the *Database*. Each object in the input file must include its "type": this field then tells us how to interpret the remaining fields, which are all assumed to be *Fact*s in the relevant Haskell type declaration.

While that sounds simple enough, recall that this is Haskell, a statically typed language. We cannot normally read in the name of a type as a string and then proceed to parse and produce an object of that type. To do so, we need a run-time description of the type (including its name) and an ability to use this description to produce an actual value of that type. Typically, this "type description" can be encapsulated in an instance dictionary, perhaps of the *FromJSON* class. Yet recall that we get the actual choice of type at *run-time*, so typical patterns of instance constraints will not suffice.

Instead, in order to parse and produce an EIGER *Database*, we must be able to look up instances at run-time. Despite sounding exotic, it turns out that this is relatively straightforward. Here is the definition:[4]

```
type Instances :: ∀ k. (k → Constraint) → Type
newtype Instances c = MkInstances (Map String (EDict c))

type EDict :: ∀ k. (k → Constraint) → Type
data EDict c where
  PackDict :: ∀ x c. c x ⇒ EDict c
```

An *Instances c* maps from the name of a type τ to an instance $c\ \tau$, as packed in an *EDict*. The *EDict* type captures an instance as a constraint on a data constructor, hiding the concrete type *x* of the argument to the class *c*. The type system here does not ensure that the key in a mapping is the string representation of the type packed in the value *EDict*; we will have to check for this when we retrieve instances from this database.

We can create a singleton *Instances* with *instanceForTypeRep*:

```
instanceForTypeRep :: ∀ x c. c x ⇒ TypeRep x → Instances c
instanceForTypeRep tr =
  MkInstances (singleton (show tr) (PackDict @x))
```

From there, because *Instances* is a *Monoid*, we can make an arbitrary collection of instances.

Retrieval is a little more awkward, because Haskell does not permit a function to return a class constraint:

```
withInstance :: Instances c → String →
  (∀ x. c x ⇒ Proxy x → r) → Maybe r
withInstance (MkInstances mapping) type_name f =
```

---

[4] In the real implementation, we have a further challenge: should type names be qualified or not? And if they are qualified, should the qualification include the package or just the modules? The implementation handles these concerns by generalizing the interface presented here to be polymorphic over the choice of type naming discipline. For simplicity in this text, we assume types are printed unqualified.

```
lookup type_name mapping ⟨&⟩
    λ(PackDict @x) → f (Proxy @x)
-- ⟨&⟩ is flip fmap, used because lookup returns a Maybe
```

The *withInstance* function must take a continuation that describes what to do with the retrieved instance. Furthermore, it must use a *Proxy* to pass the retrieved type *x*, as GHC does not yet support binding type variables in higher-rank contexts [Eisenberg 2022]. Note that the function *f* requires a *c x* constraint; this constraint is available because of the pattern-match on *PackDict*, which stored the constraint evidence.

To make *Instances* usable in practice, though, we cannot expect users to write out many invocations of *instanceForTypeRep* and combine the results. Instead, we provide the following Template Haskell function:

allGroundInstances :: *Q Type* → *Q Exp*

```
-- example usage:
$allGroundInstances [t| FromJSON |]
```

The *allGroundInstances* function looks for all instances of the requested class – *FromJSON* in the example – and creates an *Instances* that stores all of them. By "ground" here, we mean that the instance must not contain any variables. For Eiger records, which are not polymorphic, this is a reasonable restriction. Neither *allGroundInstances* nor any of the *Instances* machinery works with polymorphic instances, though it would not be hard to imagine extending them to do so, with the help of an algorithm for matching template types against concrete ones. This extra feature was not necessary in the Eiger context and so remains as future work.

Powered by *Instances*, the serializer and deserializer have the usual complexities of real-world software. Both are implemented through the use of generic programming and the generics-sop library [de Vries and Löh 2014], looking up instances via several *Instances* databases as they go.

7.3 The interface for authors

All of the frantic paddling under the surface supports a simple interface for authors (in the sense of Section 3), who we envision to be intermediate Haskell programmers working alongside a domain expert in the regulation being encoded. The last stage of our tour through Eiger highlights the design choices that make for an easier user interface.

7.3.1 Template Haskell.
In order for our high-power implementation to work, the Eiger record types must implement several different classes. This set starts with *Generic* (from *GHC.Generics*) but also includes several from generic-sop and beyond. In order to simplify the process of declaring an Eiger record, programmers must write only **deriving** *IsRecord* after their declarations (using the DeriveAnyClass extension). (The *IsRecord* class has no methods; it serves

only as a marker to be used in this mechanism.) Then, separately, users write $*deriveRecordInstances*, which looks for all *IsRecord* instances to find the set of Eiger record types and then derives all the other instances necessary. Using Template Haskell to write these instances also creates a nice abstraction, because it means that if Eiger needs new instances in the future, all we need to update is *deriveRecordInstances*; users can remain blissfully unaware of this.

7.3.2 Rule.
The *Rule* class instantiated in Section 4 is defined like this:

```
class Rule record field_name where
    rule :: HasRules record ⇒ Key record →
        OptionalRule (FactType record field_name)
type family HasRules record :: Constraint
data OptionalRule t = NoRule
                    | HasRule (ComputationRule t)
```

As we have seen, authors write *Rule* instances to define the rules for *Optional* facts. Note that *rule* returns an *OptionalRule*, meaning that users might decide not to provide a rule for a given fact; this would be used when a dataset might omit a fact, but there is still no rule to compute the missing fact from other facts.

The *rule* function also has a constraint *HasRules record*. This constraint, defined via a type family, includes *Rule* constraints for all *Optional* facts in an Eiger record type. For example, if we have

```
data Example = MkEx { ex1, ex2 :: Fact Optional Integer }
    deriving IsRecord
```

then *HasRules Example* would expand to (*Rule Example* "ex1", *Rule Example* "ex2"). The purpose of *HasRules* is to make sure that the body of the rule is able to refer to all other fields in the same record. Without *HasRules*, authors would have to ensure that all *Rule* instances for an Eiger record type are in the same module; with *HasRules*, they can be distributed across modules.[5]

Another possible design is to denote whether or not a fact has an associated rule in its type, instead of through the *NoRule*/*HasRule* constructors. There does not seem to be a concrete advantage to doing so, but as we gain experience with Eiger, it is possible this design will evolve.

7.3.3 OverloadedRecordDot.
We wish authors to be able to use familiar dot-projection syntax, like *self.tax_rate*. GHC's recent OverloadedRecordDot [Mitchell and Fletcher 2020] extension permits this: an expression *r.f* expands to a call to *getField* @"f" *r*. However, the built-in interpretation of *getField* (§7.2.2) is not powerful enough to work with

---

[5]This treatment does not extend to fields of other record types, whose *Rule* instances must be in scope in order to access the fields. It is possible to expand the *HasRules* mechanism to cover this pattern if necessary.

*ComputationRule*s. Accordingly, we use `OverloadedRecordDot` with `RebindableSyntax`.

We must now provide our own *getField* function. As *getField* must work at many different types, it must be a class method. The definition of this class follows:

type *GetField* :: *Symbol* → *Type* → *Constraint*
class *GetField field_name record* **where**
    type *FieldAccessResult field_name record* :: *Type*
    getField :: *record* → *FieldAccessResult field_name record*

Eiger includes instances to look names up in both *Key record* and *ComputationRule* (*Key record*) to flexibly allow users to project both from raw keys (such as the *self* in a *Rule* definition) and rules to produce keys (such as an expression like *self.jurisdiction*, when *self* is a *Key Entity*). Writing these (elided) instances poses no challenge.

7.3.4 DSL. To allow authors to work at the level of abstraction they understand best, the Eiger library includes a module *Eiger.DSL* that exports functions that work exclusively over *ComputationRule*s. For example, it exports the following:

(==) :: *Eq a* ⇒ *ComputationRule a* → *ComputationRule a* →
    ComputationRule Bool

    -- Used by the **if** …**then** …**else** construct
ifThenElse :: *ComputationRule Bool* → *ComputationRule a* →
    ComputationRule a → *ComputationRule a*

    -- inject ordinary values into a *ComputationRule*
use :: *a* → *ComputationRule a*
use = *pure*

In addition to these, numerical instances are provided for *ComputationRule*, so two *ComputationRule*s can be combined by +, for example. This makes for a smooth user interface, where *everything* is a *ComputationRule*.

Conclusion. We have now toured the highlights of the Eiger implementation. One of its distinctive qualities is how it works hard under the surface, using lots of fancy types and footwork, but provides a simple interface to users, putting the user experience above all.

8 Related Work

There has been a trend over the past decade of bringing digitalization and automation to the domain of legislation and regulation. The French Public Finances Directorate has been pioneering this space by encoding the French Tax Code in a DSL and publishing parts of it for the purpose of explainability [Merigoux et al. 2021b]. The OECD has been running a Rules-as-Code program that "proposes to create a machine-consumable version of some types of government rules, to exist alongside the existing natural language counterpart" [Mohun and Roberts 2020]. And the European Commission

has recently had a project on machine-readable and executable reporting requirements in the financial services industry [DG-FISMA 2020].

Several new pieces of technology have been introduced in this field. Catala is a new, text-based DSL [Merigoux et al. 2021a], while Blawx is a new, graphical DSL [Morris 2020]. These efforts are bringing natural languages and programming languages closer, coming from the direction of the latter. In contrast, so-called controlled natural languages are closing this gap coming from the other direction by restricting the grammar and vocabulary of natural languages in order to eliminate ambiguity. RegelSpraak [Fokkenrood 2011] is an implementation of this idea pursued by the Dutch Tax and Customs Administration (DTCA). Blurring the line between natural and programming language is the Agile Law Execution Factory, also developed by the DTCA [Dutch Tax and Customs Administration 2018]. It is based on JetBrain's Meta Programming System[6], a framework and IDE for building new DSLs.

These new technologies are exciting to see, and they are solving relevant problems. They also require considerable investment in the development of the ecosystem necessary for mainstream, commercial use. We are thus pursuing the approach of an embedded DSL, giving us a head-start towards commercial adoption. We are also explicit about the methodology that we believe will work in practice, which is pairing engineers with domain experts, as opposed to expecting domain experts to undergo training to write high quality code.

Regarding the run-time of Eiger, there are similarities with build systems [e.g., Mokhov et al. 2018], self-adjusting computing [Acar et al. 2007], as well as graph processing [Kaler et al. 2014; Malewicz et al. 2010]. As we scale, we will have to address performance bottlenecks as they appear. But for now it seems that the computations performed by Eiger rules are so trivial that we do not expect to be able to amortize the overhead of tracking dependencies and changes.

9 Future Work

We are currently planning to use Eiger to engage clients, putting the ideas in this paper into practice, along the lines of our outline in Section 3. We expect to use Eiger both internally for our own analysis and embedded in SaaS products for our clients. We expect Eiger to be applicable and beneficial to all the tax and legal services we offer today. Indeed, the enthusiasm for Eiger within PwC Switzerland is such that we wonder whether Eiger – or a tool inspired by it – will become a "killer app" for Haskell.

Longer term, we envision working with authorities to have them draft and publish authoritative Eiger implementations of new regulations. We believe this is where the biggest opportunities of this technology lie for society as a whole.

---

[6]https://www.jetbrains.com/mps/

On a technical level, there is still much we would like to build. For example, in the VAT domain, the typical client has several millions transactions per reporting period that need to be processed, so performance work will be a focus. The numerical hierarchy we have introduced allows us to, say, multiply a *Percentage* and a *EuroAmount* (while preventing us from multiplying two *EuroAmount*s), but does it scale beyond that? We will learn by experience. Along similar lines, we have included a few easy consultant-facing interactions, but we recognize more will be necessary for a full-fledged product. Nevertheless, our experience with Eiger as it exists – and the ease with which we have already extended it from its core – suggests that growth will be smooth.

In the end, we expect usage of Eiger to have a marked impact on PwC Switzerland. Our commitment to pair-programming has already led to successes, and the type discipline imposed by working in Haskell has helped us to understand regulation more deeply. We look forward to deploying Eiger and using it to deliver value to our clients.

References

Umut A. Acar, Matthias Blume, and Jacob Donham. 2007. A Consistent Semantics of Self-adjusting Computation. In *Programming Languages and Systems*, Rocco De Nicola (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 458–474.

Marc Andreessen. 2011. Why Software Is Eating The World. (August 2011). https://a16z.com/2011/08/20/why-software-is-eating-the-world/

Frederick P. Brooks. 1986. No Silver Bullet – Essence and Accidents of Software Engineering. In *Proceedings of the IFIP Tenth World Computing Conference*. 1069–1076.

Manuel M. T. Chakravarty, Gabriele Keller, and Simon Peyon Jones. 2005a. Associated Type Synonyms. In *International Conference on Functional Programming* (Tallinn, Estonia) *(ICFP '05)*. ACM.

Manuel M. T. Chakravarty, Gabriele Keller, Simon Peyton Jones, and Simon Marlow. 2005b. Associated Types with Class. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*.

James Cook and Oleg Grenrus. 2019–2021. some: Existential type. Hackage package. https://hackage.haskell.org/package/some

James Cook and Obsidian Systems. 2011–2020. dependent-map: Dependent finite maps (partial dependent products). Hackage package. https://hackage.haskell.org/package/dependent-map

Edsko de Vries and Andres Löh. 2014. True Sums of Products. In *Proceedings of the 10th ACM SIGPLAN Workshop on Generic Programming* (Gothenburg, Sweden) *(WGP '14)*. Association for Computing Machinery, New York, NY, USA, 83–94. https://doi.org/10.1145/2633628.2633634

DG-FISMA. 2020. Machine Readable and Executable Reporting Requirements Proof of Concept. (2020). https://etendering.ted.europa.eu/cft/cft-display.html?cftId=6051

Dutch Tax and Customs Administration. 2018. Open Rules: Agile Law Execution Factory. (2018). https://open-regels.nl/en/methoden/ALEF/

Richard Eisenberg. 2022. Modern Scoped Type Variables. GHC Proposal #448. https://github.com/ghc-proposals/ghc-proposals/pull/448 Under consideration.

Richard A. Eisenberg, Dimitrios Vytiniotis, Simon Peyton Jones, and Stephanie Weirich. 2014. Closed Type Families with Overlapping Equations. In *Principles of Programming Languages* (San Diego, California, USA) *(POPL '14)*. ACM.

Frans Fokkenrood. 2011. RegelSpraak for Business Rules: Experiences in building a Business Rules Compiler for the Dutch Tax Administration. *Business Rules Journal* 12, 11 (November 2011).

http://www.brcommunity.com/a2011/b622.html

Adam Gundry. 2017. Overloaded Record Fields. GHC Proposal #23. https://github.com/ghc-proposals/ghc-proposals/blob/master/proposals/0023-overloade Accepted.

Susanne Hammond and Mike Cowan. 2021. *Cost of Compliance 2021: Shaping the future.* Technical Report. https://legal.thomsonreuters.com/content/dam/ewp-m/documents/legal/en/pdf/reports

Tim Kaler, William Hasenplaugh, Tao B. Schardl, and Charles E. Leiserson. 2014. Executing Dynamic Data-Graph Computations Deterministically Using Chromatic Scheduling. In *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures* (Prague, Czech Republic) *(SPAA '14)*. Association for Computing Machinery, New York, NY, USA, 154–165. https://doi.org/10.1145/2612669.2612673

Csongor Kiss, Matthew Pickering, and Nicolas Wu. 2018. Generic Deriving of Generic Traversals. *Proc. ACM Program. Lang.* 2, ICFP, Article 85 (September 2018), 30 pages. https://doi.org/10.1145/3236780

José Pedro Magalhães, Atze Dijkstra, Johan Jeuring, and Andres Löh. 2010. A generic deriving mechanism for Haskell. In *Proceedings of the 3rd ACM Haskell Symposium on Haskell (Haskell '10)*. ACM.

Grzegorz Malewicz, Matthew H. Austern, Aart J.C Bik, James C. Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. 2010. Pregel: A System for Large-Scale Graph Processing. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data* (Indianapolis, Indiana, USA) *(SIGMOD '10)*. Association for Computing Machinery, New York, NY, USA, 135–146. https://doi.org/10.1145/1807167.1807184

Simon Marlow, Louis Brandy, Jonathan Coens, and Jon Purdy. 2014. There is No Fork: An Abstraction for Efficient, Concurrent, and Concise Data Access. In *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming* (Gothenburg, Sweden) *(ICFP '14)*. Association for Computing Machinery, New York, NY, USA, 325–337. https://doi.org/10.1145/2628136.2628144

Simon Marlow, Simon Peyton Jones, Edward Kmett, and Andrey Mokhov. 2016. Desugaring Haskell's Do-Notation into Applicative Operations. In *Proceedings of the 9th International Symposium on Haskell* (Nara, Japan) *(Haskell 2016)*. Association for Computing Machinery, New York, NY, USA, 92–104. https://doi.org/10.1145/2976002.2976007

Denis Merigoux, Nicolas Chataing, and Jonathan Protzenko. 2021a. Catala: A Programming Language for the Law. *Proc. ACM Program. Lang.* 5, ICFP, Article 77 (Aug. 2021), 29 pages. https://doi.org/10.1145/3473582

Denis Merigoux, Raphaël Monat, and Jonathan Protzenko. 2021b. A Modern Compiler for the French Tax Code. In *Proceedings of the 30th ACM SIGPLAN International Conference on Compiler Construction* (Virtual, Republic of Korea) *(CC 2021)*. Association for Computing Machinery, New York, NY, USA, 71–82. https://doi.org/10.1145/3446804.3446850

Neil Mitchell. 2019. Add setField to HasField. GHC Proposal #158. https://github.com/ghc-proposals/ghc-proposals/blob/master/proposals/0158-record-se Accepted.

Neil Mitchell and Shayne Fletcher. 2020. Record Dot Syntax. GHC Proposal #282. https://github.com/ghc-proposals/ghc-proposals/blob/master/proposals/0282-record-dc Accepted.

James Mohun and Alex Roberts. 2020. Cracking the code: Rulemaking for humans and machines. *OECD Working Papers on Public Governance* 42 (October 2020). https://www.oecd-ilibrary.org/governance/cracking-the-code_3afe6ba5-en

Andrey Mokhov, Neil Mitchell, and Simon Peyton Jones. 2018. Build Systems à La Carte. *Proc. ACM Program. Lang.* 2, ICFP, Article 79 (September 2018), 29 pages. https://doi.org/10.1145/3236774

Jason Morris. 2020. Blawx: Rules as Code Demonstration. *MIT Computational Law Report* (14 August 2020). https://law.mit.edu/pub/blawxrulesascodedemonstration

OECD. 2021. Tax Challenges Arising from the Digitalisation of the Economy Global Anti-Base Erosion Model Rules (Pillar Two). (December 2021).

https://www.oecd.org/tax/beps/tax-challenges-arising-from-the-digitalisation-of-the-

Simon Peyton Jones, Stephanie Weirich, Richard A. Eisenberg, and Dimitrios Vytiniotis. 2016. A reflection on types. In *A list of successes that can change the world.* Springer. A festschrift in honor of Phil Wadler.

PricewaterhouseCoopers. 2022. Machine-readable and executable regulation. (2022).

Christian R. Ulbrich, Christoph Schärer, and Stuart Jones. 2019. *What happens when the taxman gets superpowers?: A guide to the digital world of tax* (2nd ed.). PricewaterhouseCoopers AG.

A Feedback of the User Study

Here we share the verbatim responses of the user study presented in this paper. Please refer to Section 6 for the related questions.

A.1 BEPS

1. "After the second session, I felt comfortable with the approach as we were also able to cover more rules. I think it is an experience that one can get used to after a short amount of time. Also as the code focuses mainly on the mathematical approach, an in-depth knowledge of the underlying material was not necessary. Also having a small group helped to feel comfortable with the approach."

2. "It generally takes one or two rules to get into the flow during a session. Once in the flow, the sessions are really productive. Thus, I found it helpful when the session were a little longer."

3. "I am confident that we implemented our faithful interpretation of the legal text. As it is a complex matter that is still subject to a lot of interpretation and clarification, as well as outstanding guidance, the code most probably will become subject to change. We have also focused on the less complex rules thus far."

4. "Now that I'm familiar with the structure, I feel that I would be able to read someone else's implementation after been given a short introduction."

5. "Currently, I would still use the legal text as it is nicer to present. Depending on how the final tool will look, I could consider showing the code, especially when it comes to the calculation part which is shown in a more intuitive way."

6. "For me, as a personal investment, I found it really helpful to understand the rules from a more holistic and mathematical angle. Connections between rules are also becoming more obvious. I can also see the benefits of a code that can be provided to clients to perform their own calculations. [...]"

7. "I think the code is really helpful when it comes to the mathematical application. Informative text in the regulation is not reflected in the code, which still needs to be communicated to the client in another way."

8. "So far the process has been really informative for my benefit. Moreover, it is easy to work with [the engineer]. I always had time to think through a rule or change it and I never felt rushed or incompetent. In my point of view, to structure and organize the project/process more, I would suggest to have a timeline/going to market strategy including milestones prepared in order for everyone to have a clear goal/deadline insight."

A.2 CBAM

1. "The step-by-step, article-by-article approach makes it easy to follow, to understand, and ultimately to implement the coding."

2. "Seeing the result of the code straight away would be interesting. If, for instance, some dummy data was used, what would the result of the code be? Would it result in an error, in an incorrect result, or the correct result?"

3. "In my view, the CBAM text is relatively simple – articles are relatively short and there are not many cross-references to articles. I have high confidence that the legal text is faithfully implemented."

4. "Reading and understanding other people's code can be a challenge. Assuming, however, that other people follow a logical and structured way of coding, I have high confidence that I could read and understand their implementation of legal text."

5. "My preference would be to use the code but also include the legal text as non-executable comments to provide context."

6. "The pair-programming approach is essential, particularly in the beginning. Coders need to get an understanding of the legal text's logic, and subject matter experts (SMEs) need to understand the logic of the code. The benefits goes both ways. I actually do believe that SMEs benefit most. Dissecting the legal text in code-able parts may also help SMEs to make implicit assumptions explicit, and to look at the legal text from a different perspective."

7. "No, I think the code could replicate the legal text without struggle."

8. "My overall feedback is very positive. The process of transforming legal text to code is clear and easy to follow. One challenge that SMEs may struggle with is the direction of logic. Legal text often starts with a general claim, followed by specific conditions or exceptions. Code often works the other way around. Specifics are first defined, and more general claims are build upon this. I therefore think it is important to include non-executable comments in the code to provide context."

A.3 UK Value Added Tax Regime

1. "I found it much easier to get to grips with the approach than I had expected. I had expected it to take some time for us to build the ground rules for a specific domain so that we could create something that

worked, but I was surprised that we were already working on functional code within 5 minutes of starting. This is one of the key benefits of pairing two people with specific domain expertise (one "software technical" and one "domain technical") so each can focus on their specialism and speed is therefore increased. I was surprised at how quickly I could pick up the logic of the code that was being written with my input."

2. "If it is possible to identify some terms we can either infer or have as default values so that the creation of rules or facts can be streamlined to a degree, that would help make the experience better."

3. "I am confident that what we built would faithfully represent the basic functionality of the VAT system. There are a few edge cases that would specifically need to be incorporated for it to be fully representative of the framework, but I am confident that with a bit more coding time we would be able to represent the full set of scenarios."

4. "One of the benefits of the pair-programming approach is that you experience the code being built "first hand", meaning that you understand why each term is being used. Reviewing it "second hand" means you do not have some of this context, but having said that, the code was surprisingly easy to read even for examples where I had not been present at writing. I would say certainly no more difficult than reading formulas in Excel, just different, and not beyond the capability of a domain expert to learn within a reasonably short timeframe."

5. "If not the code, what is missing for you to use this instead? When advising on regulation, we must default to the text that is legally enforceable. At this point, this is the legal text – particularly as the industry is built on human language. However, with faithfully represented code like the examples here, I could see the industry moving to a place where the code could be used to advise on implementation. It would be important to annotate the code with references and links to legal texts to aid the review of the implemented code and "bridge the gap" between the "old" and the "new" systems."

6. "Absolutely"

7. "Not for the primary activities of regulation: aggregation of data, comparing to a simple rule book to derive conclusions, and comparatively simple calculations. The code can deliver on these with very low execution risk in my view. The main challenge will be the replication of areas of judgement; situations where option A or B could be a possibility, and the "right" choice is not necessarily clear, or is a tactical/strategic one. Indeed, it is these questions that makes one organization's risk appetite different from others, and is the key reason for having a regulatory compliance

team. We need the ability with Eiger to accommodate these decision points, making the options clear and then letting a user make a decision."

8. "I really enjoyed the session, was pleasantly surprised with the speed of establishing executable code from a blank page. I really think this framework has potential, and the process is effective at "upskilling" both developers on the workings of regulation, and domain experts on how to think like a developer."