

Special Section on Quantitative Systems Pharmacology: A Foundation to Establish Precision Medicine

Optimizing IFN Alpha Therapy against Myeloproliferative Neoplasms^S

Gurvan Hermange, Paul-Henry Cournède, and Isabelle Plo

Université Paris-Saclay, CentraleSupélec, Laboratory of Mathematics and Informatics (MICS), Gif-sur-Yvette, France (G.H., P.-H.C.); INSERM U1287, Villejuif, France (I.P.); Gustave Roussy, Villejuif, France (I.P.); and Université Paris-Saclay, Villejuif, France (I.P.)

Received December 15, 2022; accepted June 21, 2023

ABSTRACT

Myeloproliferative neoplasms (MPNs) are hematologic malignancies that result from acquired driver mutations in hematopoietic stem cells (HSCs), causing overproduction of blood cells and an increased risk of thrombohemorrhagic events. The most common MPN driver mutation affects the *JAK2* gene (*JAK2*^{V617F}). Interferon alpha (IFN α) is a promising treatment against MPNs by inducing a hematologic response and molecular remission for some patients. Mathematical models have been proposed to describe how IFN α targets mutated HSCs, indicating that a minimal dose is necessary for long-term remission. This study aims to determine a personalized treatment strategy. First, we show the capacity of an existing model to predict cell dynamics for new patients from data that can be easily obtained in clinic. Then, we study different treatment scenarios in silico for three patients, considering potential IFN α dose-toxicity relations. We assess when the treatment should be interrupted depending on the response, the patient's age, and the inferred development of the malignant clone without IFN α . We find that an optimal strategy would be to treat patients with a constant

dose so that treatment could be interrupted as quickly as possible. Higher doses result in earlier discontinuation but also higher toxicity. Without knowledge of the dose-toxicity relationship, trade-off strategies can be found for each patient. A compromise strategy is to treat patients with medium doses (60–120 μ g/week) for 10–15 years. Altogether, this work demonstrates how a mathematical model calibrated from real data can help build a clinical decision-support tool to optimize long-term IFN α therapy for MPN patients.

SIGNIFICANCE STATEMENT

Myeloproliferative neoplasms (MPNs) are chronic blood cancers. Interferon alpha (IFN α) is a promising treatment with the potential to induce a molecular response by targeting mutated hematopoietic stem cells. MPN patients are treated over several years, and there is a lack of knowledge concerning the posology strategy and the best timing for interrupting therapy. The study opens avenues for rationalizing how to treat MPN patients with IFN α over several years, promoting a more personalized approach to treatment.

Introduction

BCR-ABL–negative myeloproliferative neoplasms (MPNs) are hematologic malignancies that include essential thrombocythemia, polycythemia vera, and primary myelofibrosis leading to overproduction of myeloid cells (red cells, platelets, and

granulocytes). These diseases are clonal due to the expansion over decades of a hematopoietic stem cell (HSC) that has undergone the acquisition of a genetic abnormality (Van Egeren et al., 2021; Hermange et al., 2022; Williams et al., 2022). After homologous recombination, homozygous malignant clones can develop in parallel to heterozygous clones. The mutations affect *JAK2*, calreticulin (*CALR*), and the thrombopoietin receptor (*MPL*) genes, with *JAK2*^{V617F} being the most prevalent in more than 95% of polycythemia vera and around half of essential thrombocythemia and primary myelofibrosis. All these mutations are gain-of-function leading to constitutive activation of the *JAK2*/*STAT* pathway (Vainchenker et al., 2011).

Thrombohemorrhagic complications are the main comorbidities even if the transformation into secondary leukemia of dismal prognosis occurs in a significant fraction of cases (between

This work is supported in part by a grant from the Prism project and by the Agence Nationale de la Recherche [Grant ANR-18-IBHU-0002] (to P.-H.C. and I.P.). This work was also supported by grants from INCA Plbio2018, 2021 (to I.P.) and Ligue Nationale Contre le Cancer (équipe labellisée 2019) and by the Appel à Projets Pré-néoplasies 2021 (C21021LS) with financial support from Institut Thématique Multi-Organismes (ITMO) Cancer of Aviesan within the framework of the 2021–2030 Cancer Control Strategy, with funds administered by INSERM.

The authors declare no competing financial interests.
dx.doi.org/10.1124/jpet.122.001561.

^S This article has supplemental material available at jpet.aspetjournals.org.

ABBREVIATIONS: CF, clonal fraction; HMR, high molecular response; HSC, hematopoietic stem cell; IFN, interferon; MPN, myeloproliferative neoplasm; MSE, mean square error; VAF, variant allele frequency.

5% and 20%). Clinicians generally treat patients with MPN to normalize the blood parameters and/or to improve symptoms, mainly using aspirin, cytoreductive treatment (hydroxyurea), and JAK1/2 inhibitors [ruxolitinib and fedratinib (Inrebic)]. However, neither of these treatments impacts the malignant clone selectively, and the only curative option is bone marrow allogeneic transplantation in the most severe myelofibrotic patients. In this context, pegylated interferon alpha 2a (Peg-IFN α 2a) is a promising treatment. When used, IFN α is mainly a second-line therapy. It harbors high rates of hematologic responses in $JAK2^{V617F}$ MPN patients and some molecular responses, including high molecular responses (HMRs), even after treatment arrest (Kiladjian et al., 2008; Gisslinger et al., 2020; Mascarenhas et al., 2022). In clinic, the molecular response refers to measuring the variant allele frequency (VAF) in mature cells and quantifying how the VAF decreases over the therapy. Despite these positive effects, some $JAK2^{V617F}$ MPN patients do not respond to IFN α , and HMR is only observed in $\sim 20\%$ of $JAK2^{V617F}$ patients. Long-term treatment of 2–5 years is required to obtain HMR. Additionally, toxicity increases with the IFN α dose (Yamane et al., 2008), resulting in the choice of dose de-escalation in many patients after 1 year, potentially jeopardizing the success of long-term HMR (Mosca et al., 2021), all of which results in decision-making challenges for clinicians. Therefore, this therapy could be improved and optimized to increase the molecular response while limiting toxicity.

Previously, it has been shown that the IFN α targets the $JAK2^{V617F}$ HSCs in a preclinical mouse model (Hasan et al., 2013; Mullally et al., 2013; Austin et al., 2020). Due to the difficult access to human HSCs by biologic methods, the effect of IFN α on the mutated human HSC dynamics was investigated using a tailor-made compartmental mathematical model calibrated with data from a prospective, observational, and longitudinal cohort of treated patients. Our objective with this model was to quantify not only the molecular response among mature cells but also the response at the stem cell level. We inferred that the IFN α slowly targets human $JAK2^{V617F}$ HSCs and that the higher the dose, the better the response (Mosca et al., 2021). The latter mathematical model was then extended to account for the dose variations during the treatment when the mutated HSCs were targeted (Hermange et al., 2021). This model is used in this present work to develop a clinical decision-support tool for predicting the response of MPN patients to IFN α and optimizing their long-term

therapy. First, we verify that we can get accurate predictions from the model, clinical observations from a new patient, and prior knowledge from a cohort of patients. Then, we explore several therapeutical strategies to find which one would minimize the IFN α toxicity until the end of the treatment. Finally, combining the previous mathematical model with another one describing how $JAK2^{V617F}$ -mutated HSCs—when not treated—expand over time (Hermange et al., 2022), we assess when the therapy could be interrupted depending on how the patient responds to IFN α , their age, and the proportion of homozygous and heterozygous $JAK2^{V617F}$ HSCs (i.e., the zygosity).

Materials and Methods

Experimental and Clinical Data

The experimental and clinical data used in this study are those of 19 $JAK2^{V617F}$ patients from Mosca et al. (2021). For each patient of this cohort, we have:


- Their age at the beginning of the treatment.
- Measurements of the VAF in mature cells (in granulocytes) at different time points during the treatment. We define them as clinical data since these observations are molecular biology data obtained during routine clinical exams.
- Observations of the clonal architecture [i.e., measurement of the heterozygous and homozygous clonal fractions (CF)] at different time points during the treatment. We define them as experimental data since they are obtained in biology research laboratories.
- The variations of the dose they received over time.

We display the experimental and clinical data of patient #32 and those of each patient (Fig. 1; Supplemental Fig. A, 1–19). We denote by \mathcal{D} the dataset for all 19 patients (see Supplemental Methods A: Data and observation model).

Predicting the Long-Term Dynamics

The model studied in Mosca et al. (2021) (see Fig. 2; Supplemental Methods B: Model) was calibrated from observations on 19 patients with MPN having the mutation $JAK2^{V617F}$. A hierarchical Bayesian inference method was used to estimate the parameters of each individual in addition to a population effect, thus reducing the risk of overfitting. Then, the minimal IFN α dose $d_{\text{inf}}^{(i)}$ required for a patient i to reach long-term remission was estimated (Supplemental Methods C2.4). Our objective is now to demonstrate how the model could be used as a clinical decision-support tool. In clinical routines, trimestrial measurements of the clonal architecture of progenitor cells are not

Fig. 1. Longitudinal observations of patient #32. The x-axis corresponds to the time from the start of the IFN α therapy. Clinical data consist of measurements of the VAF in mature cells (black squares); experimental data consist of measurements of the heterozygous (green circles) and homozygous (blue triangles) CF in progenitor cells. The shaded beige areas correspond to the weekly dose of IFN α received over time.



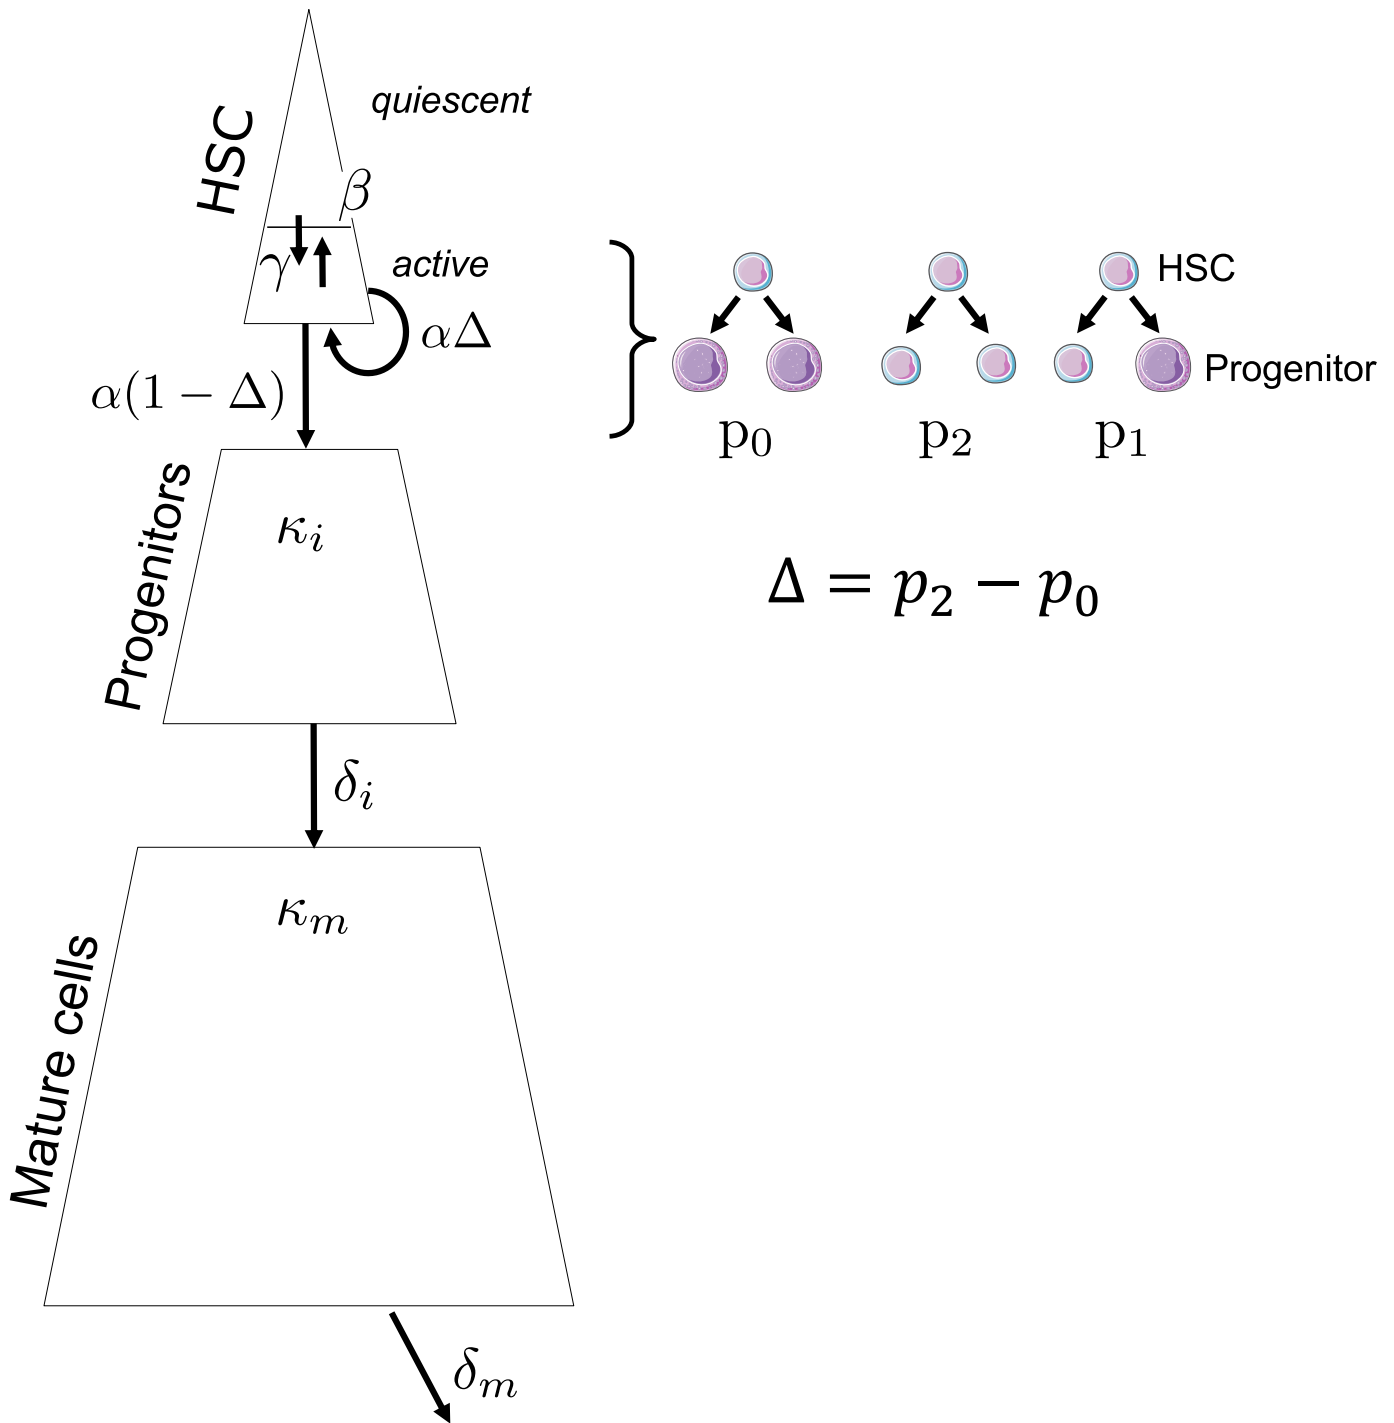


Fig. 2. Schematic representation of the model for one cell population (e.g., mutated $JAK2^{V617F}$ heterozygous cells). Quiescent HSCs can exit quiescence at a rate $\gamma N_q(t)$ with $N_q(t)$ as the number of quiescent HSCs at time t . Active HSCs can become quiescent at a rate $\beta N_a(t)$ or be recruited to divide at a rate $\alpha N_a(t)$ with $N_a(t)$ as the number of active HSCs at time t . In the latter case, an HSC could divide symmetrically (producing two HSCs) with probability p_2 , asymmetrically with probability p_1 , or produce two immature cells (progenitors) with probability p_0 . $\Delta = p_2 - p_0$ corresponds to the balance between symmetrical and differentiated division. Progenitor cells expand with a rate constant κ_i and become mature cells with a rate constant δ_i . Mature cells expand with a rate constant κ_m and die at a death rate δ_m .

feasible. However, the clonal architecture of heterozygous and homozygous mutated immature cells (i.e., the measurement of CF) yields essential information that determines the response to the treatment as shown in Tong et al. (2021). In terms of cost, time, and human resources, it can be considered reasonable to measure the CF twice for each patient: one just before they start their therapy and one afterward (generally after about 300 days of therapy; we explore the impact of this choice in the Supplemental Methods F: Estimating the

best timing for measuring the clonal architecture). The VAF in mature cells can easily be measured during routine clinical follow-up.

If the model were to be used by physicians, then we would have to evaluate its capacity to predict cell dynamics for new patients for whom such data can be easily obtained in clinical routines. Yet, we do not have access to patients outside the cohort of Mosca et al. (2021) to study the predictive capabilities of the model and then to study different dose strategies. Therefore, we choose to consider our 19 $JAK2^{V617F}$

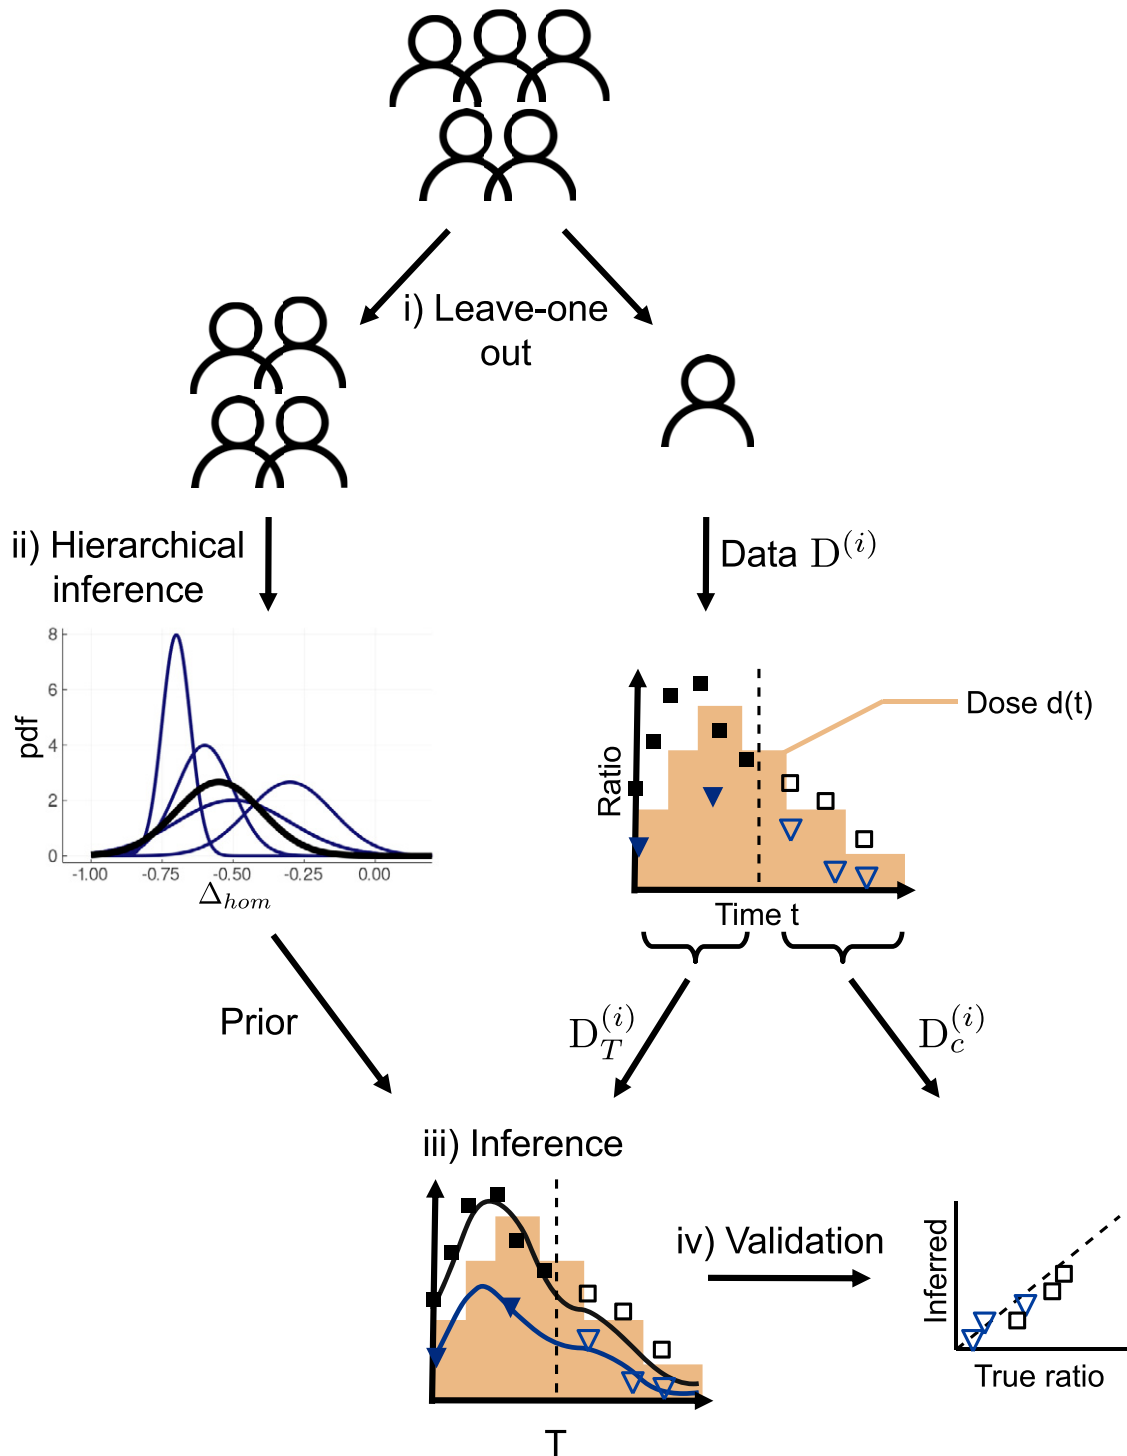


Fig. 3. Illustration of the method for predicting the long-term cell dynamics over IFN α therapy. Removing one patient from the cohort (i), we run a hierarchical Bayesian inference procedure using the data of the remaining patients (ii). That is, for each model parameter (here, for example, Δ_{hom}), we estimate its probability density (pdf) for each patient (blue lines) as well as for the patients' population (black line). The latter will be retained as prior knowledge. The data $D^{(i)}$ of the excluded patient from the initial cohort consists of VAF measurements in mature cells (black squares) and CF in progenitor cells, here homozygous CF (blue triangles), over time t . We also have the information on the dose of IFN α $d(t)$ that the patient received (in brown). We split their dataset into two parts, the first one $D_T^{(i)}$ corresponding to observations before a time T (that we call assimilation time) and the second one $D_c^{(i)}$ to the observations after that time. Based on the dataset $D_T^{(i)}$ and the prior obtained from the remaining cohort of patients, we infer the dynamics of the mutated cells on treatment (iii). The black line corresponds to the dynamics of the VAF and the blue one to the dynamics of the homozygous CF. To evaluate the capacity of the model to predict the long-term dynamics correctly, we compare the model predictions and the control dataset $D_c^{(i)}$ that was not used for estimating the model parameters.

TABLE 1
List of the main notations

Notation	Meaning
\mathcal{D}	Dataset of the 19 considered $JAK2^{V617F}$ patients
i	Index of the patient
t	Time
T	Assimilation time (until which observations are considered for calibrating the model)
τ	Patient's age at which the treatment could be interrupted
$\mathcal{D}^{(i)}$	Dataset of patient i
\mathcal{D}_{-i}	Data of all patients except patient i
$\mathcal{D}_T^{(i)}$	Observations of patient i , before time T , used for calibrating the model
$\mathcal{D}_c^{(i)}$	Observations of patient i not used for the model calibration
d	Normalized dose of IFN α (between 0 and 1; 1 corresponding to 180 $\mu\text{g}/\text{week}$)
$z(d)$	IFN α toxicity (as a function of the dose d)
$M(\tau)$	Amount of IFN α administered over $[T, \tau]$ (when penalizing for high dose)

-mutated patients with MPN, from which we will remove one individual (leave-one-out) and for whom we will perform data assimilation (Fig. 3i; Supplemental Methods C: Parameter estimation).

We denote $\mathcal{D}_{-i} = \mathcal{D} \setminus \mathcal{D}^{(i)}$ as the dataset for all patients except individual i (Table 1 summarizes the main notations). In this article, we will consider three individuals who were each in turn excluded from the cohort: patients #12, #18, and #32. These three patients are chosen because several observations of their response before $t = 300$ days of treatment are available, which is only the case for a few patients in the cohort (see Supplemental Methods A). Furthermore, each of the three patients exhibits a distinctive type of response to IFN α , making their study interesting: patient #12 has mostly homozygous cells, and they seem to be slowly but steadily targeted over the treatment; patient #18 presents both homozygous and heterozygous mutated cells initially, in the same proportions, and only the homozygous clone is targeted during the treatment; and patient #32 responds with a so-called ‘‘bell curve,’’ meaning that their VAF and CF first strongly increase at the beginning of the treatment and then decrease.

We run three hierarchical estimation procedures (Supplemental Methods C1) corresponding respectively to the datasets \mathcal{D}_{-12} , \mathcal{D}_{-18} , and \mathcal{D}_{-32} . Of these three calibrations, we will retain only the estimate of the population distribution (Fig. 3ii), which will be used as prior distribution to predict the response of the left-apart patient i . The progressive inclusion of observations for this patient (through what is called in Bayesian statistics the likelihood) will allow the updating of the previous distribution (i.e., the prior). We refer to this statistical process as ‘‘data assimilation’’ in this article (Fig. 3iii; Supplemental Methods C2). Let $T \in \{300, 600, 1000\}$ [days] be the times (of assimilation) until which we consider the observations of patient i . The dataset of patient i used for the model calibration will be denoted by: $\mathcal{D}_T^{(i)}$. Note that this dataset includes all VAF measurements before time T and only two measurements of the clonal architecture. The other observations of the CF among progenitor cells will not be used for estimating the model parameters. Indeed, in clinical routines, we cannot have repeated measurements of the progenitor clonal architecture. Thus, we choose to use only some of the observations available for progenitor cells before the assimilation time to treat the patient's data as if they were obtained from realistic processes in clinical routine. We consider the observation of the progenitor CF at the start of the therapy and about 300 days afterward.

For patient i , for the assimilation time $T \in \{300, 600, 1000\}$, we will then estimate the *posterior* distribution of each parameter from the prior distribution and the observations $\mathcal{D}_T^{(i)}$. The parameters will then be used to infer the on-treatment dynamics of the VAF in the mature cells as well as the heterozygous and homozygous CF in the progenitors and the HSCs (Fig. 3iii).

The confrontation of the predicted values with the observations not used for the model calibration will allow us to evaluate the quality of the predictions (Fig. 3iv). The control dataset for assessing the quality of the predictions will be denoted as:

$$\mathcal{D}_c^{(i)} = \mathcal{D}^{(i)} \setminus \mathcal{D}_T^{(i)} \quad (1)$$

This control dataset includes all observations made after time T as well as the observations of the CF among progenitors that were not used for estimating the model parameters.

To quantify how well the model fits the observations, either those used to estimate the model parameters ($\mathcal{D}_T^{(i)}$) or those used as control ($\mathcal{D}_c^{(i)}$), we will compute the mean square error (MSE) (see Supplemental Methods C2.5). In particular, we will check how $\text{MSE}_{\text{pred}}^{\text{VAF}}$, which confronts the predicted VAF values to the ones truly observed after the assimilation time T (i.e., not used for the model calibration), evolves for higher T values.

Optimizing the Therapy

Once we have evaluated the ability of the model to predict the long-term response from data potentially available in clinical routines, we can explore different therapeutic strategies (see *Therapeutic Strategies*) and find the optimal one (Fig. 4). The optimality will be defined as the minimal dose of IFN α administered until treatment interruption (see *Assessing When to Interrupt the Treatment*) when penalizing higher doses associated with higher risks of toxicity (see *Minimizing Drug-Related Toxicity during the Treatment*). The complete method is presented in Supplemental Methods E: Optimizing the therapy, whereas the focus here is on the most important aspects for understanding the article.

Assessing When to Interrupt the Treatment. IFN α therapy against MPN is a long-term therapy. It has been proven to induce, in some cases, an HMR, which is an undetectable level of $JAK2^{V617F}$ -mutated cells in peripheral blood (Kiladjian et al., 2008). However, we cannot clinically assess whether all mutated HSCs have been completely eradicated and whether a relapse might occur in the future. $JAK2^{V617F}$ -positive MPNs are hematologic malignancies that develop over decades before the onset of the symptoms, as shown recently in different studies (Van Egeren et al., 2021; Hermange et al., 2022; Williams et al., 2022). It was shown that heterozygous (het) $JAK2^{V617F}$ -mutated HSCs, when escaping stochastic extinction, would expand on average according to the following dynamics (Hermange et al., 2022):

$$\tilde{N}_{\text{het}}(t) = \tilde{N}_{0,\text{het}} \exp(s_{\text{het}}(t - t_0)) \quad (2)$$

with a fitness estimated to be equal to $s_{\text{het}} = 20.4\%$ per year and $\tilde{N}_{\text{het}}(t)$ corresponding to the total number of mutated heterozygous HSCs over time t during the MPN development in the absence of IFN α therapy.

The dynamic of the homozygous clonal expansion was not studied in Hermange et al. (2022). Here, we generalize previous findings and assume that the homozygous clone would grow exponentially like the heterozygous one but with a higher fitness $s_{\text{hom}} > s_{\text{het}}$, as shown by Williams et al. (2022). We estimate $s_{\text{hom}} = 1.21 \times s_{\text{het}}$ (Supplemental Methods D3).

It is further assumed that the clonal expansion given by eq. 2, inferred from individuals before they received IFN α , will also be valid after they interrupt the treatment. Then, when inferring the CF of heterozygous and homozygous mutated HSCs from the model, at any time during the




Fig. 4. Illustration of the method for optimizing the IFN α therapy. Considering the observations for a patient until time T and their model parameters estimated, we can obtain the dynamics of the mutated cells (here, for example, the CF of homozygous progenitors in blue) over time t for a given dose $d(t)$. We can then consider different therapeutic strategies [i.e., different ways $d(t)$ will evolve for $t \geq T$] (i). Here, for example,

therapy, we could estimate using eq. 2 what would be the future development of the malignant clones after a treatment interruption at time τ . We can thus deduce the time $t > \tau$ at which the VAF among HSCs would exceed 7.5%. This value corresponds to the classic risk threshold above which it is considered that there could be a risk of thrombosis or cardiovascular events (i.e., a reappearance of the disease) (Dupont et al., 2007). The later this time, the better for the patient. Ideally, the therapy should be stopped at a time τ so that the relapse does not occur during the patient's life. In practice, considering a life expectancy at 65 years old to be equal to 90 years old (Kontis et al., 2017), we will consider interrupting the treatment of a patient at an age τ such that the VAF among HSCs stays below 7.5% when the patient is under 90.

Thus, the criterion to interrupt the treatment depends on the age of the patients, the way they respond to the treatment, and their zygosity.

In this study, mutations associated with MPN such as *DNMT3A*, *TET2*, and *TP53*, which could potentially undergo a clonal expansion independent of the *JAK2*^{V617F} clonal dynamics, are not considered.

In the Supplemental Methods D: Interrupting the therapy, we detail the choice of the criterion for the treatment interruption (Supplemental Methods D1), the model used for describing the clonal expansion (Supplemental Methods D2), and the estimation of the homozygous *JAK2*^{V617F} fitness (Supplemental Methods D3).

Therapeutic Strategies. If we can correctly estimate the model parameters of a new patient (and thus correctly predict their mutated cell dynamics when they receive the dose $d(t)$ for $t \geq T$, with T as the assimilation time), we could explore alternative therapeutic strategies to those currently used by the clinicians. In practice, clinicians generally increase the dose until they observe a hematologic response and then de-escalate the dose, as was observed in the cohort of Mosca et al. (2021). However, this strategy might not be optimal since it can lead to early relapse.

In this article, different treatment strategies are studied, all of which will involve some parameters to optimize (see *Minimizing Drug-Related Toxicity during the Treatment*). Then, the strategies can be compared and the best one can be chosen (according to a criterion that we define later).

The list of strategies studied in this article is far from exhaustive, and we stick to simple ones that clinicians could easily employ. It is considered that the strategy will be set up from the assimilation time T to the interruption time τ . Before time T , the clinician might have treated the patient with a dose escalation, allowing an assessment of the patient's tolerance to the treatment. During the period from the start of the therapy to T , measurements are made and will allow the estimation of the model parameters.

The three different treatment strategies studied are the following (Fig. 5):

- Constant:

$$d(t) = \bar{d} \text{ for } t \in [T, \tau] \quad (3)$$

- Periodic:

$$d(t) = \begin{cases} \bar{d} & \text{for } t \in [T + 2kL, T + (2k + 1)L], k \in \mathbb{N} \\ d_{inf} & \text{for } t \in [T + (2k + 1)L, T + (2k + 2)L], k \in \mathbb{N} \end{cases} \quad (4)$$

we illustrate a therapeutic strategy that is periodic. We can estimate the dynamics of the mutated cells for this new posology and assess the time τ when the CF will be low enough so that, after treatment interruption, the further clonal expansion will not lead to the onset of MPN symptoms before the patient is 90 years old (ii). Note that the illustration is done considering progenitor cells but that in our calculations, we estimate τ based on the inferred CF among HSCs. We then consider one scenario for the increase of IFN α toxicity z with the dose d (here, convex) (iii). Finally, we can compute $M(t)$ for $T \leq t \leq \tau$, which describes the amount of IFN α administered over the time interval $[T, \tau]$ when penalizing for high doses (iv). The quantity of interest is $M(\tau)$, which we want to minimize across different therapeutic strategies and for which we want to evaluate the impact of the dose-toxicity relation.

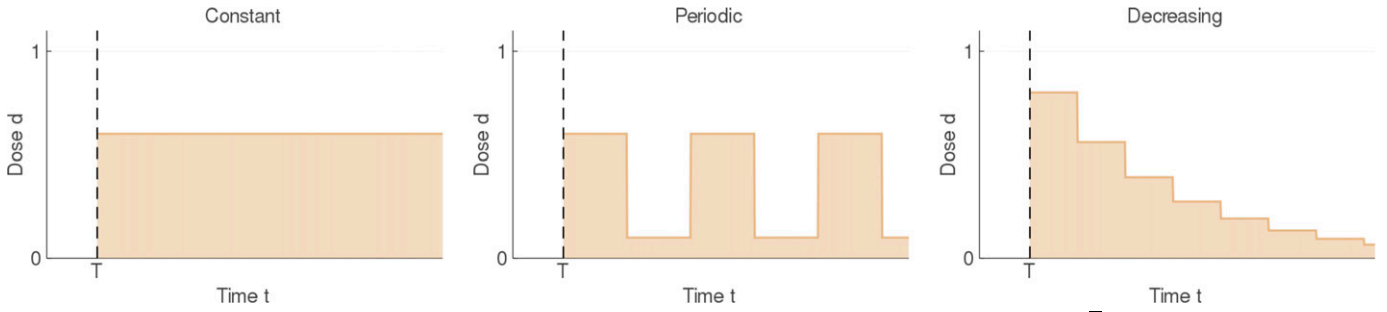


Fig. 5. Illustration of the three different therapeutic strategies that we consider. On the left, a constant dose over time (here, $\bar{d}=0.6$). In the middle, a periodic strategy (here, $\bar{d}=0.6$, $L = 4$ months, and d_{inf} assumed to be equal to 0.1). On the right, a decreasing strategy (here, $\bar{d}=0.8$, $L = 3$ months, and $\lambda = 0.7$).

- Decreasing:

$$d(t) = \begin{cases} \bar{d} & \text{for } t \in [T, T + L], \\ \lambda d(t - L) & \text{for } t \in [T + L, \tau] \end{cases} \quad (5)$$

The first strategy involves only one parameter \bar{d} to optimize, the second strategy two parameters (\bar{d} and the period L), and the last strategy three parameters, namely \bar{d} , λ , and the period L . This third strategy is the one often implemented in practice. Note that in the equation above, the expression of the decreasing function d is recursive (see Supplemental Methods E3 for another way to express the decreasing strategy).

d_{inf} corresponds to the minimal personalized dose that can be computed from the estimated parameter vector (Supplemental Methods C2.4). For all strategies, we consider $\bar{d} \geq d_{inf}$, $\lambda \in [0, 1]$. We also want $L \geq 3$ months to avoid too-frequent changes of posologies (regarding the treatment duration of generally many years). Besides, for too-high values of L , the decreasing and periodic strategies are equivalent to the constant one. To avoid finding ourselves in that case, the following condition is set: $L \leq 2$ years. In addition, λ is constrained to be inferior to 0.95 to avoid the case where the decreasing strategy would approach the constant one when $\lambda \rightarrow 1$. The optimal parameters will be estimated by doing a grid search (Supplemental Methods E1).

To note, we voluntarily choose not to consider stop-and-go strategies. Indeed, patients might need some time to tolerate the treatment; therefore, it should not be stopped to be resumed later to avoid this adaptation phase unless, of course, the interruption is permanent.

The periodic strategy, as defined by imposing the minimal dose to be equal to d_{inf} , is then the strategy that is the closest to the stop-and-go strategy.

Minimizing Drug-Related Toxicity during the Treatment. The main objective of the study is to optimize the IFN α therapy against MPN. For that purpose, several potential therapeutic strategies were considered in *Therapeutic Strategies*. We should define the clinically relevant quantity that has to be optimized. Intuitively, we want the treatment to be as short as possible. With the considered model, the corresponding optimization problem would be solved by choosing to give the maximal IFN α dose of 180 $\mu\text{g}/\text{week}$. However, the toxicity of IFN α is known to increase with the dose (Yamane et al., 2008), so optimizing only the duration of treatment would result in high-toxicity strategies. Thus, a penalty has to be applied for high doses. Toxicity includes everything harmful to the patient, either hematologic or not. In particular, IFN α is known to be associated with side effects such as depression (Trask et al., 2000; Lotrich et al., 2007), flu symptoms, or thrombocytopenia. Since no data exist to quantify the dose-toxicity relation, we will consider several hypothetical scenarios and evaluate their impact on the results. We define $z(d)$ the drug toxicity as a function of the dose and consider four potential behaviors for the dose-toxicity relation (see Fig. 6):

- Linear:
 $z(d) = 2d$
- Convex:

$$z(d) = 3d^2$$

- Concave:

$$z(d) = \frac{3}{2}\sqrt{d} \quad (6)$$

- Composite:

$$z(d) = \frac{5\sqrt{5}}{9}\sqrt{d-0.1} \mathbf{1}_{d \in [0.1, 0.55]} + \frac{5}{4\sqrt{5}} \frac{1}{\sqrt{1-d}} \mathbf{1}_{d \in [0.55, 1]}$$

For normalization, the average toxicity is set to 1 for the four formulations:

$$\int_0^1 z(d) dd = 1 \quad (7)$$

For the composite relation, we assume that there is no toxicity below a low threshold of 0.1 (corresponding to 18 $\mu\text{g}/\text{week}$), then a concave behavior for $d \in [0.1, 0.55]$ followed by a convex one where the toxicity tends to infinity for $d \rightarrow 1$. The value delimiting both behaviors (i.e., $d = 0.55$) is chosen so that the derivative of z is continuous over $[0.1, 1]$. More details concerning the construction of the composite relation are provided in Supplemental Methods E4.

$M(t)$ is defined as the toxicity-related amount of IFN α administered over $[T, t]$:

$$M(t) = \int_T^t z(d(u)) du \quad (8)$$

where d (as a function of time) is assumed to be either constant, periodic, or decreasing and z (as a function of d) is either linear, convex,




Fig. 6. The four different scenarios that we consider for modeling the way the IFN α toxicity may increase with the dose.

concave, or composite. Then, the quantity to minimize is $M(\tau)$. In the linear scenario, $M(\tau)$ would then correspond to the total amount of IFN α administered between T and τ . To note that, in eq. 8 we do not consider the interval $[0, T]$ since we only want to optimize the treatment from $t = T$.

We study for patients #12, #18, and #32 three therapeutic strategies and four scenarios of dose toxicity (i.e., twelve different conditions). For each of them, we will estimate the parameters related to the therapeutic strategy that minimizes the quantity $M(\tau)$ (Supplemental Methods E1). The strategies are then compared based on this value, for a given hypothetical scenario of toxicity, to find the optimal one. Moreover, in the absence of prior knowledge on how IFN α toxicity increases with the dose, a trade-off strategy can be proposed (Supplemental Methods E2).

Results

Minimal Observations Are Sufficient To Predict the Long-Term Mutated Cell Dynamics

First, we investigated whether some measurements of the VAF and only two measurements of the progenitor clonal architecture were sufficient to predict the long-term response to the treatment. To present the results in this section, we will focus on patient #18 (all of our results are detailed in Supplemental Material G: Detailed results; the results for patients #12 and #32 concerning the prediction part are detailed in Supplemental Material G1). The cell dynamics observed during IFN α therapy for patient #18 are interesting since they initially presented both heterozygous and homozygous clones (and in the same proportions), but the homozygous clone was targeted during the treatment when the heterozygous one continued to expand. Such dynamics might be challenging to predict when the parameter estimation is mainly based on VAF measurements.

For patient #18, we assess whether we could predict the evolution of their VAF correctly, but also both their heterozygous and homozygous CF in progenitors (Fig. 7), from:

- only two measurements of the clonal architecture in immature cells (at the start of the treatment and 248 days afterward),
- several measurements of the VAF before a time $T \in \{300, 600, 1000\}$ [days], and
- prior knowledge obtained from the remaining patients.

By progressively adding more information on the VAF dynamics (i.e., increasing the assimilation time), we increase the confidence in our predictions. It is particularly true for the VAF and the homozygous CF. As early as $T = 300$ days, when only four VAF measurements are used for the parameter estimation, we observe a good agreement between the predicted values, both for the VAF and the CF, and the experimental values not used for estimating the model parameters.

MSE_{pred}^{VAF} —quantifying the error between the observed and predicted VAF for different assimilation times T —is equal to:

- $6.6 \cdot 10^{-3}$ for $T = 300$ (10 VAF measurements used for computing MSE_{pred}^{VAF})
- $5.8 \cdot 10^{-4}$ for $T = 600$ (six measurements used for computing this MSE)
- $5.2 \cdot 10^{-4}$ for $T = 1000$ (three measurements used for computing this MSE)

We show that adding more observations for the VAF improves the quality of our prediction. However, the (median) inferred dynamics for the mutated heterozygous progenitors are closer to the observations for $T = 600$ days than $T = 1000$, resulting for the latter in a higher MSE on the heterozygous CF measurements not used for the inference (Supplemental Table G2). For $T = 1000$, we overestimate the heterozygous CF, illustrating the difficulty in correctly predicting both the dynamics of heterozygous and homozygous cells when the inference is mainly based on the VAF measurements.

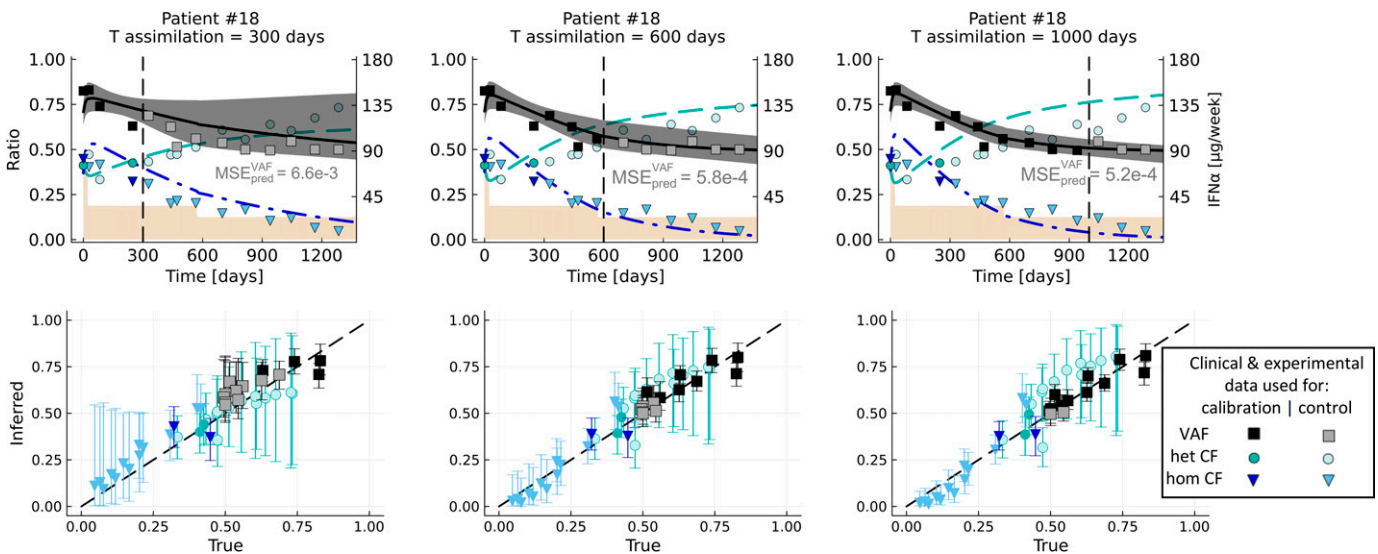


Fig. 7. Data assimilation results for patient #18. From left to right, the assimilation times are $T = 300, 600,$ and 1000 days (vertical dash line). At the top, we present the predicted dynamics. In black are shown the inferred VAF (median value and 95% credibility interval), and in green and blue are shown the inferred CF (median values) of heterozygous and homozygous mutated progenitors, respectively. The squares correspond to the clinical data and the circles and triangles to the experimental data, either used to calibrate the model (dark colors, $\mathcal{D}_T^{(i)}$) or to control the quality of the predictions (light colors, control dataset $\mathcal{D}_c^{(i)}$). We represent the IFN α dose variations during treatment (in $\mu\text{g}/\text{week}$) in brown. The MSE displayed just below the VAF dynamics is computed on the VAF measurements not used for the calibration (i.e., the VAF measurements after the assimilation time T) such that MSE_{pred}^{VAF} quantifies the predicted error on the VAF. At the bottom, we compare the inferred values (median values, on the y-axis) with those observed (on the x-axis). The error bars correspond to the 95% credibility intervals.

For patient #12, we also observe a good agreement between the predicted VAF and the observed ones, with an MSE_{pred}^{VAF} of $1.1 \cdot 10^{-3}$ and $1.5 \cdot 10^{-3}$ for T equal to 300 and 600, respectively. However, we systematically underestimate the true homozygous CF (Supplemental Fig. G1; Supplemental Material G1.1; Supplemental Table G1). For patient #32, we obtained good predictions, both for the VAF and the homozygous CF (patient #32 has almost no heterozygous clones) for $T = 300$ or 1000 days (with MSE_{pred}^{VAF} respectively equal to $2.17 \cdot 10^{-2}$ and $2.35 \cdot 10^{-2}$) but, surprisingly, poor predictions for $T = 600$ days ($MSE_{pred}^{VAF} = 0.25$) (Supplemental Fig. G5; Supplemental Material G1.3; Supplemental Table G3). But, if instead of considering the measurement of the clonal architecture at time $t = 392$, we chose to take the observation at $t = 508$, we observed that the predictions for $T = 600$ days were far better ($MSE_{pred}^{VAF} = 7.5 \cdot 10^{-3}$) (Supplemental Fig. G6; Supplemental Material G1.4; Supplemental Table G4), illustrating the importance of the experimental design to make accurate predictions. We explore this question in Supplemental Material G3.

From the results on the three patients considered in this study, we conclude that it is possible to make predictions sensible of the mutated cell dynamics from VAF observations and two measurements of the CF among progenitor cells but that the choice of the time point for the clonal architecture matters, raising the concern of how the observation time should be chosen. After 300 days of treatment, it is possible to get good predictions of the cell dynamics and after 600 days to have good confidence in our predictions.

The Dose Should Be Kept Constant until the Treatment Is Stopped

We studied whether other therapeutic strategies could lead to better results (i.e., faster molecular response while limiting the potential IFN α toxicity). To evaluate these strategies, we focused on patient #32 (results for patients #12 and #18 are detailed in Supplemental Material G2). For this patient, the model predicts that a decrease in the dose after 2000 days of therapy might slow down the molecular response and even induce a risk of relapse (Supplemental Fig. G6). The way the patient was treated, with a dose escalation over the first 3 years of therapy up to the maximal dose of 180 $\mu\text{g}/\text{week}$, then a de-escalation, is also characteristic of the strategies applied in clinical routine and already observed in Mosca et al. (2021).

We consider $T = 600$ days. Using the observations before that time (five VAF observations, two clonal architectures: one at $t = 0$, the second one at $t = 508$) and the prior knowledge obtained from the 18 remaining patients, we estimated the posterior distribution of the model parameters and showed that we obtain good predictions (Supplemental Material G1.4).

We now consider that the actual parameter vector is the estimated one (mean of the posterior distribution) and study how different dose strategies impact the response of the treatment after $T = 600$ days. With the estimated parameter vector, the minimal dose (under which the treatment would not sufficiently target the mutated HSCs, resulting in a relapse) is estimated to be $d_{inf} = 0.345$ (i.e., 62 $\mu\text{g}/\text{week}$). We consider different scenarios for how the drug toxicity increases as a function of the dose, as presented in Fig. 6. For each scenario, we study the three therapeutic strategies presented in *Therapeutic Strategies*, for which we find the parameters (e.g., the choice of the dose \bar{d}) that minimize the value of $M(\tau)$. This latter value corresponds to the toxicity-related

amount of IFN α administered from T to the time τ when the therapy could be interrupted (see *Assessing When to Interrupt the Treatment*). We can compare the three therapeutic strategies for a given drug-toxicity relation and select the one with the best (i.e., the lowest) value for $M(\tau)$. Results are presented in Fig. 8. It turns out that the constant strategy is the optimal one among the four studied scenarios, also for patients #12 (Supplemental Fig. G8) and #18 (Supplemental Fig. G9). With the two other strategies, namely the periodic one and the decreasing one, we obtain at the optimum lower values of $M(\tau)/(\tau - T)$ (that we interpret as values of the yearly toxicity-related amount of IFN α) compared with the constant strategy. However, having periods with low doses or decreasing the dose over time also delays treatment interruption. Therefore, these two strategies are less effective in terms of treatment duration.

However, the choice of the optimal dose highly depends on how the IFN α toxicity increases with the dose. If there were a sharp increase in the toxicity for low doses followed by a slower one (concave scenario), we would recommend using high doses (about 180 $\mu\text{g}/\text{week}$ for patient #32). Indeed, only slightly decreasing the dose would not strongly impact the toxicity, whereas decreasing the dose to a too-large extent would delay the treatment interruption or even induce a relapse if the dose is below a minimum value, as shown in Hermange et al., (2021). If, on the contrary, the toxicity would only sharply increase for high doses (convex scenario), then we would recommend using medium doses (about 90 $\mu\text{g}/\text{week}$ for patient #32). Eventually, for the linear dose-toxicity relation (which is both concave and convex), or for the composite relation (which is first concave, then convex), we find that the recommended dose for patient #32 would be about 135 $\mu\text{g}/\text{week}$. Given the lack of existing data quantifying the relationship between the dose and the toxicity of IFN α , we cannot conclude which dose level would be optimal. Our results highlight the need to estimate such relations since we demonstrate their importance in deciphering how patients with MPN should be treated.

In the absence of prior knowledge on whether the dose toxicity would better be described by a linear, concave, convex, or composite relationship, we can determine the strategy that, even if not optimal, would be the best compromise. Finding a trade-off strategy is important since the best strategy under the hypothesis that the dose-toxicity relation would be concave would be detrimental under another hypothesis. We illustrate that point with patient #32. In the concave scenario, the best strategy is to treat patient #32 at a constant dose of 180 $\mu\text{g}/\text{week}$ over about 5 years. This value would actually correspond to the highest dose considered in the study (i.e., a very high dose). If the dose-toxicity relation were actually not concave but linear, such a strategy would still be acceptable, with a value for $M(\tau)$ equal to 4020, placing it in the top 3.4% and better than the best periodic strategy. But, in the case where the dose-toxicity relationship was convex, treating patient #32 with such a high dose would be harmful. Indeed, it turns out that the optimal strategy in the concave scenario is also the worst one in the convex scenario.

Therefore, without having assessed in the first 2 years of the therapy how the patient responds to the treatment or in the absence of data that would quantify the dose-toxicity relation for IFN α , the best choice would be to select the trade-off strategy.

This strategy is the one that gives good results under the four hypothetical dose-toxicity relationships. For patient #32, the trade-

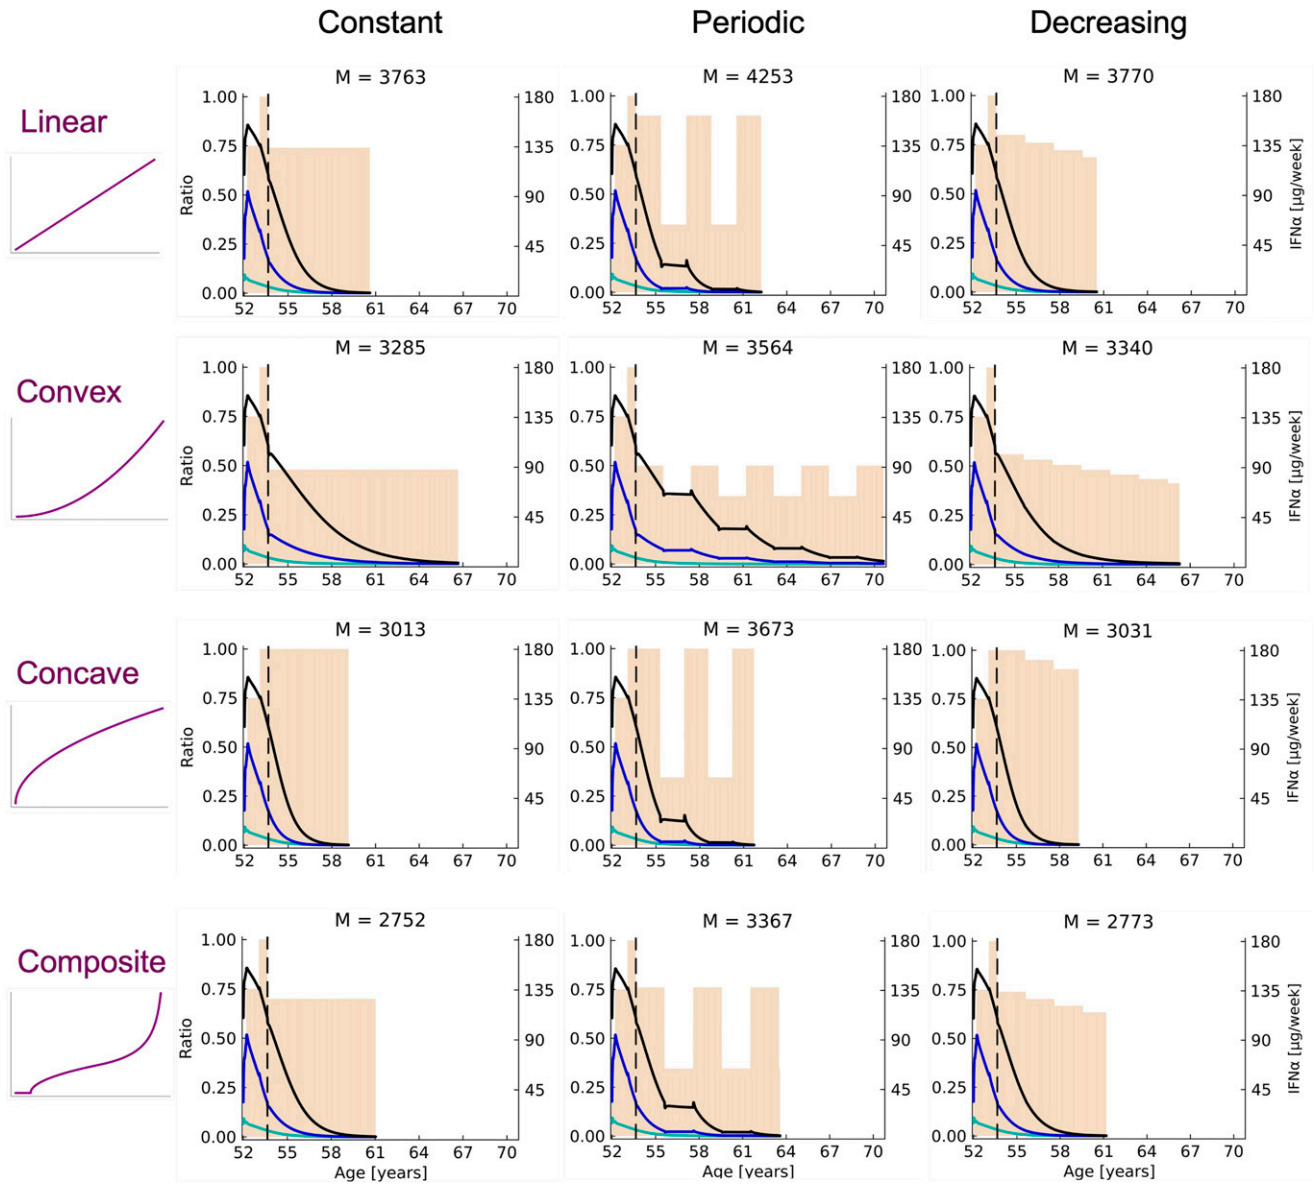


Fig. 8. Results of the treatment optimization for patient #32. Four hypothetical relationships between IFN α toxicity and the dose are studied: linear (at the top), convex (on the second line), concave (on the third line), and composite (at the bottom). For each of them, three therapeutic strategies—constant on the left, periodic on the middle, and decreasing on the right—are optimized to minimize $M(\tau)$ (denoted by M on the figures), which is the toxicity-related amount of IFN α administrated from T (vertical dashed line) to the treatment interruption time τ (when the weekly dose of interferon, represented in beige, drops to zero). The dynamics of the VAF (black line), the heterozygous CF (green line), and the homozygous CF (blue line) will differ for $t \geq T$ according to how the doses vary (depending on the considered therapeutic strategy and its parameter values). M values are only comparable for a given dose-toxicity relation—the lower, the better. If the actual dose-toxicity relation were linear, the constant strategy would be optimal for $\bar{d} = 135 \mu\text{g}/\text{week}$ (with the associated age of treatment interruption $\tau = 61$ years and an optimal value of $M(\tau) = 3763$); the periodic strategy would be optimal for $\bar{d} = 153 \mu\text{g}/\text{week}$ and $L = 22$ months; and the decreasing strategy would be optimal for $\bar{d} = 144 \mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best. We can see how the optimal parameters of the decreasing strategy actually tend to make it approach a constant strategy, which would be the limiting case when $\lambda \rightarrow 1$ and/or $L \rightarrow \infty$. It is precisely to avoid finding oneself in this limiting case that constraints were imposed on the parameters (Supplemental Methods E1). If the actual dose-toxicity relation were convex, the constant strategy would be optimal for $\bar{d} = 86 \mu\text{g}/\text{week}$; the periodic strategy would be optimal for $\bar{d} = 90 \mu\text{g}/\text{week}$ and $L = 23$ months; and the decreasing strategy would be optimal for $\bar{d} = 100 \mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best. If the actual dose-toxicity relation were concave, the constant strategy would be optimal for $\bar{d} = 180 \mu\text{g}/\text{week}$; the periodic strategy would be optimal for $\bar{d} = 180 \mu\text{g}/\text{week}$ and $L = 20$ months; and the decreasing strategy would be optimal for $\bar{d} = 180 \mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best. If the actual dose-toxicity relation were composite, the constant strategy would be optimal for $\bar{d} = 126 \mu\text{g}/\text{week}$; the periodic strategy would be optimal for $\bar{d} = 137 \mu\text{g}/\text{week}$ and $L = 2$ years; and the decreasing strategy would be optimal for $\bar{d} = 133 \mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best.

off would be to treat him with a constant dose of $115 \mu\text{g}/\text{week}$ over 8 years until the treatment could be interrupted, here at age 62 (i.e., after about 10 years of IFN α therapy since patient #32 began the treatment at age 52). This trade-off strategy would be:

- in the top 1.7% if the dose-toxicity relation were linear,
- in the top 2.3% if the dose-toxicity relation were convex,

- in the top 2.9% if the dose-toxicity relation were concave, and
- in the top 0.2% if the dose-toxicity relation were composite.

For patient #18, the trade-off strategy would also be the constant one, with $\bar{d} = 61 \mu\text{g}/\text{week}$, until treatment discontinuation at the age of 72 (i.e., after about 15 years

of IFN α therapy). For patient #12, the trade-off strategy would be the decreasing one, with $\bar{d} = 86 \mu\text{g}/\text{week}$, $\lambda = 0.45$, and $L = 16$ months, until treatment discontinuation at the age of 77 (i.e., after about 14 years of IFN α therapy) (Supplemental Material G2).

Discussion

We proposed a mathematical approach combining modeling, statistical inference, and optimization to rationalize the manner in which we treat $JAK2^{V617F}$ patients with MPN in the long term using IFN α . Clinicians often proceed to a dose de-escalation when they observe a hematologic response. The rationale behind this strategy is that IFN α is associated with general or hematologic side effects such as depression (Trask et al., 2000; Lotrich et al., 2007), flu symptoms, or thrombocytopenia and that the toxicity increases with the dose (Yamane et al., 2008). In the absence of both data quantifying the dose-toxicity relation and information on the response of mutated HSCs to the treatment, their therapeutic strategy, albeit empirical, is relevant. However, as observed from the cohort of Mosca et al. (2021), clinicians sometimes have to increase the dose again after a treatment interruption or a strong decrease when a relapse is observed. To avoid such behaviors that could harm the patient, we explored which therapeutic strategy would be optimal given hypothetical dose-toxicity relations. In particular, we estimated which dose level should be administered. We found that treating the patients with a constant dose until the therapy is interrupted rather than decreasing or alternating a low and a higher dose periodically should be more relevant. Even if not considered initially as a selection criterion, the simplicity of the constant strategy involving only one value to choose (i.e., the IFN α dose to be kept constant until the treatment interruption) makes it easy to deploy in clinics. Such a simple strategy appears more optimal than more sophisticated ones. However, the number of strategies that we explored was far from exhaustive; thus, we could not exclude that patients with MPN could be treated more successfully with different strategies. In our mathematical approach, we do not guarantee that, after interrupting the treatment, there will not be a relapse. On the contrary, we quantified this risk and chose to interrupt the treatment so that the reappearance of the MPN disease would not occur before the age of 90. This age threshold is based on the estimated life expectancy of 65-year-old individuals in developed countries (Kontis et al., 2017). Increasing this age, thus reducing the risk, would result in treating the patient for a longer time. It might also seem relevant to adjust this criterion regarding the sex of the patient, given that women have a higher life expectancy than men. Concerning the reappearance of the MPN disease, we considered that the relapse might occur when the VAF in HSCs exceeds 7.5% (Dupont et al., 2007). However, some patients with essential thrombocythemia might exhibit lower VAF, whereas a higher VAF is necessary for others to get the MPN symptoms. The choice of this value particularly matters when considering the problem of early screening (Hermange et al., 2022). In our case, however, the patients who have interrupted their treatment would still be followed so that clinicians might detect quickly if and when a relapse occurs.

Given hypothetical dose-toxicity relations, we found that treating a patient with a constant dose is optimal. However, the value of the dose depends on how IFN α toxicity increases with the dose. High doses should be recommended if the

patient tolerates the treatment well; medium or low doses should be recommended otherwise. No data quantify, in the general case, what the dose-toxicity relation could be. Such a relation could actually be patient dependent. In our approach, we looked for the best therapeutic strategy once the patient had already been treated for 600 days. Therefore, we could already get prior insight into how the patient responded to a dose escalation. In other words, the clinician might provide information on the potential dose-toxicity relation, so we could recommend the constant dose value to administer to the patient. Involving the clinicians in getting such prior information would benefit our method, which could, in turn, provide more relevant guidelines on how it might be optimal to treat the patient. Without such dose-toxicity relation, we can find a trade-off strategy. This strategy is such that the dose-toxicity-related quantity of IFN α administered to the patient would be as little as possible for all four scenarios of dose-toxicity that we considered. It is worth noting that, for optimizing the therapy, we only considered minimizing a quantity related to the drug toxicity and did not account for economic criteria, for example. The issue of the economic cost of the therapy was studied by Pedersen et al. (2020). We also limit ourselves to four potential dose-toxicity relations, where the toxicity would strictly increase with the dose. It should not be excluded that more complex relations might exist and that the toxicity could, for example, reach a plateau. With the four dose-toxicity relations that we considered, we found that the trade-off strategy would also be when the dose is maintained at a constant value for two of the three patients studied. For the third patient, the decreasing strategy was found to be the best compromise; our results are patient dependent, and more patients should be studied to identify general patterns. In particular, the dose level is specific to each patient and, more precisely, to how they respond to the treatment. For the two patients whose trade-off strategy was the constant one, we found different optimal values of 61 and 115 $\mu\text{g}/\text{week}$. Our study considered any potential dose between 0 and 180 $\mu\text{g}/\text{week}$. In practice, clinicians can inject 135 or 180 μg of IFN α . By choosing the period between two medication intakes, we can have a range of potential doses, but some dose values would still not be clinically relevant. Furthermore, when the frequency of IFN α injection is low, below one every 10 days (Xu et al., 1998), the modeling assumption that the dose input is a piecewise constant function might not be justified. It would then be worth studying pharmacokinetic/pharmacodynamic (PK/PD) models [for a review see, for example, Gabrielsson and Green (2009)].

Our mathematical approach focuses on the mutated $JAK2^{V617F}$ HSCs: how IFN α targets them and then how they expand after the treatment interruption. The fact that HSCs cannot be observed in vivo justifies using mathematical models to infer their dynamics. Here, we combined two models: one that describes the dynamics of mutated cells during the therapy (Mosca et al., 2021) and one that describes the clonal expansion of the mutated cells in the absence of IFN α (Hermange et al., 2022). We assumed that the stochastic effect could be neglected. However, when the treatment lasts for a long time, the number of mutated cells might reach a low value, so deterministic models might not be appropriate anymore. It would then be relevant to propose a stochastic model for how IFN α would selectively target a few mutated HSCs within a large pool of wild-type (normal) HSCs. Concerning the wild-type HSCs, our model is based on the hypothesis that their number stays constant, even during disease development.

Alternative models could also be proposed, either for the effect of the treatment (Ottesen et al., 2020; Pedersen et al., 2021) or the clonal expansion without IFN α (Van Egeren et al., 2021; Williams et al., 2022). In particular, the model from Pedersen et al. (2021) and Ottesen et al. (2020) presents the advantage of having been initially designed for calibration from VAF measurements; thus, it might be more relevant for use in clinical routine. The counterpart is that they do not account for the zygosity, which is a determinant of the response to IFN α (Mosca et al., 2021; Tong et al., 2021). A roundabout way exists to access information on zygosity in mature cells from a VAF measurement: using the 46/1 haplotype (Hasan et al., 2014) for patients heterozygous for this polymorphism (it would also be possible to determine other polymorphisms as informative). Indeed, it has been shown that generally, in the case of mitotic recombination by which a heterozygous *JAK2*^{V617F}-mutated cell gives a homozygous mutated cell, the cell will also become homozygous for the 46/1 haplotype. Thus, measuring the VAF for this haplotype will tell us the proportion of homozygous mutated cells. Even if the model from Mosca et al. (2021) was not originally calibrated from data that could be easily obtained in clinical routine, we showed that having only two observations of the progenitor clonal architecture would be sufficient to get accurate predictions when having some prior information on the parameter distribution. This prior knowledge came from a hierarchical Bayesian inference when using all available information about the heterozygous and homozygous CF dynamics in progenitor cells (and not only two measurements of the clonal architecture) from other patients. Still, the prediction quality is uneven depending on the timing for measuring the clonal architecture. Poor predictions also mean that our recommendations for treating the patient could not be adapted. This is why we explored the issue of optimal experimental design: how to choose the best timing for measuring the clonal architecture (i.e., the one leading to the most accurate inferences). Intensive *in silico* investigations have shown that the time for the second CF measurement is important, but a search for any systematic was inconclusive (Supplemental Materials F and G3). It turned out that the best timing for a given patient might not be the best for another, preventing us from finding a general pattern of when it would be optimal to measure the CF among progenitor cells for a new patient. Thus, it might be worth adding more patients to the study. In addition, the experimental design question could be extended to the dates at which the VAF should be measured. We only superficially addressed the issue of experimental design, which deserves to be studied in more detail. However, if clinicians can easily use our model-based predictions and treatment recommendations, it would be more complicated to make recommendations on when they should collect patient blood samples. Indeed, the timing for collecting patient observations is subject to scheduling constraints over which we have no influence. It would be difficult to impose the dates when the patient should have an appointment with their clinicians. For that reason, we limited ourselves to the experimental design study and only considered the question of the best timing for the second measurement of the progenitor clonal architecture. Even if we could not conclude when it would be optimal for a new patient to get that measurement, we showed that not all choices were equal.

Finally, we demonstrated how a mathematical model calibrated from real data could predict MPN patients' response to IFN α and optimize their long-term therapy. The proposed methods were designed to be easily applied as a decision-support tool in clinical

routine. Our mathematical approach can be extended to study other drugs against MPN such as ropeg-IFN α 2b (Barbui et al., 2021; Gisslinger et al., 2020; Mascarenhas et al., 2022), which is a different pegylated IFN α from the one studied in this article, or even its combination with ruxolitinib (Kiladjian et al., 2022). Furthermore, our methods could be applied more broadly to studying other chronic hematologic malignancies.

Acknowledgments

The authors thank D. Madhavan for her help in editing the manuscript.

Data Availability

The authors declare that all of the data supporting the findings of this study are available within the paper and its Supplemental Material.

Authorship Contributions

Participated in research design: Hermange, Cournède, Plo.

Performed data analysis: Hermange, Cournède, Plo.

Wrote or contributed to the writing of the manuscript: Hermange, Cournède, Plo.

References

- Austin RJ, Straube J, Bruedigam C, Pali G, Jacquelin S, Vu T, Green J, Gräsel J, Lansink L, Cooper L, et al. (2020) Distinct effects of ruxolitinib and interferon-alpha on murine *JAK2*V617F myeloproliferative neoplasm hematopoietic stem cell populations. *Leukemia* **34**:1075–1089.
- Barbui T, Vannucchi AM, De Stefano V, Masciulli A, Carobbio A, Ferrari A, Ghirardi A, Rossi E, Ciceri F, Bonifacio M, et al. (2021) Ropoginterferon alfa-2b versus phlebotomy in low-risk patients with polycythaemia vera (Low-PV study): a multicentre, randomised phase 2 trial. *Lancet Haematol* **8**:e175–e184.
- Dupont S, Massé A, James C, Teyssandier I, Lécluse Y, Larbret F, Ugo V, Saulnier P, Koscielny S, Le Couédic JP, et al. (2007) The *JAK2* 617V>F mutation triggers erythropoietin hypersensitivity and terminal erythroid amplification in primary cells from patients with polycythemia vera. *Blood* **110**:1013–1021.
- Gabrielsson J and Green AR (2009) Quantitative pharmacology or pharmacokinetic pharmacodynamic integration should be a vital component in integrative pharmacology. *J Pharmacol Exp Ther* **331**:767–774.
- Gisslinger H, Klade C, Georgiev P, Krochmalczyk D, Gercheva-Kyuchukova L, Egyed M, Rossiev V, Dulicek P, Illes A, Pylypenko H, et al.; PROUD-PV Study Group (2020) Ropoginterferon alfa-2b versus standard therapy for polycythaemia vera (PROUD-PV and CONTINUATION-PV): a randomised, non-inferiority, phase 3 trial and its extension study. *Lancet Haematol* **7**:e196–e208.
- Hasan S, Cassinat B, Droin N, Le Couédic JP, Favale F, Monte-Mor B, Lacout C, Fontenay M, Dosquet C, Chomienne C, et al. (2014) Use of the 46/1 haplotype to model *JAK2*(V617F) clonal architecture in PV patients: clonal evolution and impact of IFN α treatment. *Leukemia* **28**:460–463.
- Hasan S, Lacout C, Marty C, Cuingnet M, Solary E, Vainchenker W, and Villeval J-L (2013) *JAK2*V617F expression in mice amplifies early hematopoietic cells and gives them a competitive advantage that is hampered by IFN α . *Blood* **122**:1464–1477.
- Hermange G, Rakotonirainy A, Bentrion M, Tisserand A, El-Khoury M, Girodon F, Marzac C, Vainchenker W, Plo I, and Cournède P-H (2022) Inferring the initiation and development of myeloproliferative neoplasms. *Proc Natl Acad Sci USA* **119**:e2120374119.
- Hermange G, Vainchenker W, Plo I, and Cournède PH (2021) Mathematical modeling, selection and hierarchical inference to determine the minimal dose in IFN alpha therapy against Myeloproliferative Neoplasms. *arXiv preprint* DOI: 10.48550/arXiv.2112.10688 [published ahead of print].
- Kiladjian J-J, Cassinat B, Chevret S, Turlure P, Cambier N, Roussel M, Bellucci S, Grandchamp B, Chomienne C, and Fenaux P (2008) Pegylated interferon-alpha-2a induces complete hematologic and molecular responses with low toxicity in polycythemia vera. *Blood* **112**:3065–3072.
- Kiladjian J-J, Ianotto J-C, Soret J, Maslah N, Chaffaut C, Barraco F, Dubruy V, Capron C, Tisserand A, Rolland-Neyret V, et al. (2022) Final results of ruxopeg, a phase 1/2 adaptive randomized trial of ruxolitinib (rux) and pegylated interferon alpha (ifna) 2a in patients with myelofibrosis (mf). *Blood* **140** (Suppl 1):577–578 DOI: 10.1182/blood-2022-156389.
- Kontis V, Bennett JE, Mathers CD, Li G, Foreman K, and Ezzati M (2017) Future life expectancy in 35 industrialised countries: projections with a Bayesian model ensemble. *Lancet* **389**:1323–1335.
- Lotrich FE, Rabinovitz M, Gironda P, and Pollock BG (2007) Depression following pegylated interferon-alpha: characteristics and vulnerability. *J Psychosom Res* **63**:131–135.
- Mascarenhas J, Kosiorek HE, Prchal JT, Rambaldi A, Berenson D, Yacoub A, Harrison CN, McMullin MF, Vannucchi AM, Ewing J, et al. (2022) A randomized phase 3 trial of interferon- α vs hydroxyurea in polycythemia vera and essential thrombocythemia. *Blood* **139**:2931–2941.
- Mosca M, Hermange G, Tisserand A, Noble R, Marzac C, Marty C, Le Sueur C, Camparion H, Vertenoil G, El-Khoury M, et al. (2021) Inferring the dynamics of mutated

- hematopoietic stem and progenitor cells induced by IFN α in myeloproliferative neoplasms. *Blood* **138**:2231–2243.
- Mullally A, Bruedigam C, Poveromo L, Heidele FH, Purdon A, Vu T, Austin R, Heckl D, Breyfogle LJ, Kuhn CP, et al. (2013) Depletion of Jak2V617F myeloproliferative neoplasm-propagating stem cells by interferon- α in a murine model of polycythemia vera. *Blood* **121**:3692–3702.
- Ottesen JT, Pedersen RK, Dam MJB, Knudsen TA, Skov V, Kjær L, and Andersen M (2020) Mathematical modeling of mpns offers understanding and decision support for personalized treatment. *Cancers (Basel)* **12**:2119.
- Pedersen RK, Andersen M, Knudsen TA, Sajid Z, Gudmand-Hoeyer J, Dam MJB, Skov V, Kjær L, Ellervik C, Larsen TS, et al. (2020) Data-driven analysis of JAK2V617F kinetics during interferon-alpha2 treatment of patients with polycythemia vera and related neoplasms. *Cancer Med* **9**:2039–2051.
- Pedersen RK, Andersen M, Knudsen TA, Skov V, Kjær L, Hasselbalch HC, and Ottesen JT (2021) Dose-dependent mathematical modeling of interferon- α -treatment for personalized treatment of myeloproliferative neoplasms. *Comput Syst Oncol* **1**:e1030 DOI: 10.1002/cso2.1030.
- Tong J, Sun T, Ma S, Zhao Y, Ju M, Gao Y, Zhu P, Tan P, Fu R, Zhang A, et al. (2021) Hematopoietic stem cell heterogeneity is linked to the initiation and therapeutic response of myeloproliferative neoplasms. *Cell Stem Cell* **28**: 502–513.e6.
- Trask PC, Esper P, Riba M, and Redman B (2000) Psychiatric side effects of interferon therapy: prevalence, proposed mechanisms, and future directions. *J Clin Oncol* **18**:2316–2326.
- Vainchenker W, Delhommeau F, Constantinescu SN, and Bernard OA (2011) New mutations and pathogenesis of myeloproliferative neoplasms. *Blood* **118**:1723–1735.
- Van Egeren D, Escabi J, Nguyen M, Liu S, Reilly CR, Patel S, Kamaz B, Kalyva M, DeAngelo DJ, Galinsky I, et al. (2021) Reconstructing the lineage histories and differentiation trajectories of individual cancer cells in myeloproliferative neoplasms. *Cell Stem Cell* **28**:514–523.e9.
- Williams N, Lee J, Mitchell E, Moore L, Baxter EJ, Hewinson J, Dawson KJ, Menzies A, Godfrey AL, Green AR, et al. (2022) Life histories of myeloproliferative neoplasms inferred from phylogenies. *Nature* **602**:162–168.
- Xu Z-X, Hoffman J, Patel I, and Joubert P (1998) Single-dose safety/tolerability and pharmacokinetic/pharmacodynamics (PK/PD) following administration of ascending subcutaneous doses of pegylated-interferon (PEG-IFN) and interferon α -2a (IFN α -2a) to healthy subjects. *Hepatology* **28** (Suppl):702A.
- Yamane A, Nakamura T, Suzuki H, Ito M, Ohnishi Y, Ikeda Y, and Miyakawa Y (2008) Interferon- α 2b-induced thrombocytopenia is caused by inhibition of platelet production but not proliferation and endomitosis in human megakaryocytes. *Blood* **112**:542–550.

Address correspondence to: Paul-Henry Cournède, Université Paris-Saclay, CentraleSupélec, 3 Rue Joliot Curie, 91190 Gif-sur-Yvette, France. E-mail: paul-henry.cournede@centralesupelec.fr; or Dr. Isabelle Plo, INSERM U1287, Gustave Roussy, PR1, 114 Rue Edouard Vaillant, 94805 Villejuif, France. E-mail: isabelle.plo@gustaveroussy.fr

Optimizing IFN alpha therapy against Myeloproliferative Neoplasms

Gurvan Hermange¹, Paul-Henry Cournède^{1*}, and Isabelle Plo^{2, 3, 4*}

¹Université Paris-Saclay, CentraleSupélec, Laboratory of Mathematics and Informatics (MICS), Gif-sur-Yvette, France.

²INSERM U1287 (INSERM, Gustave Roussy, Université Paris-Saclay), Villejuif, France

³Gustave Roussy, Villejuif, France

⁴Université Paris-Saclay, Villejuif, France

* corresponding authors: paul-henry.cournede@centralesupelec.fr ; isabelle.plo@gustaveroussy.fr

Supplemental Materials and Methods

We present the data used in this work in section A, and our model in section B. Section C presents the parameter estimation procedure that allows us to make long-term predictions of the effect of IFN α therapy against MPN. In § C.1, we present how the hierarchical Bayesian estimation is performed, and in § C.2, how we estimate the parameters for a new patient.

In section D, we describe how we model the expansion of the $JAK2^{V617F}$ heterozygous and homozygous mutated clones in the absence of IFN α , and then, the criterion for deciding at which age the therapy could be interrupted.

In sections E and F, we detail the methods used for optimizing the therapy and the experimental design, respectively.

The detail of our results is presented in section G. In § G.1, we detail the results of the prediction; in § G.2, those for the optimization; and in § G.3, those for the experimental design.

Contents

A	Data and observation model	3
B	Model	12
C	Parameter estimation	13
C.1	Hierarchical inference	13
C.2	Parameter estimation for a new patient	14
C.2.1	Prior	14
C.2.2	Bayesian Inference	14
C.2.3	Uncertainty propagation	14
C.2.4	Estimating d_{inf}	15
C.2.5	Assessing the quality of the fits and predictions	15
D	Interrupting the therapy	16
D.1	Criterion	16
D.2	Model of the clonal expansion without $IFN\alpha$	16
D.3	Estimating the fitness of the $JAK2^{V617F}$ homozygous clone	17
D.3.1	Approach	17
D.3.2	Finding the relation $f : S_{het} \mapsto s_{het}$	18
D.3.3	Results	19
E	Optimizing the therapy	20
E.1	Grid search of the optimal parameters	20
E.2	Trade-off strategy	22
E.3	Note about the decreasing strategy	22
E.4	Note about the composite dose-toxicity relation	22
F	Estimating the best timing for measuring the clonal architecture	24
G	Detailed results	27
G.1	Prediction	27
G.1.1	Patient #12	27
G.1.2	Patient #18	30
G.1.3	Patient #32	33
G.1.4	Patient #32 when choosing another time point for measuring the clonal architecture	34
G.1.5	Synthesis	37
G.2	Optimization	39
G.2.1	Patient #12	39
G.2.2	Patient #18	40
G.2.3	Patient #32	40
G.3	Experimental Design	44

A Data and observation model

The data we use in this study come from [15]. In particular, we consider the 19 $JAK2^{V617F}$ MPN patients studied in [9]. Observations originally consisted of the clonal architecture measurements of progenitors and the proportion of mutated alleles (VAF - Variant Allele Frequency) in mature cells. We denote by $\mathcal{D} = \{\mathcal{D}_i\}_{i \in \{1, \dots, N\}}$ the dataset for all $N = 19$ patients. For a patient i , the data consist of observations at different times $t_k^{(i)}$ from the beginning of the therapy ($t = 0$). The VAF $\hat{y}_k^{(i)}$ is measured in peripheral blood, that is, among mature cells. Such a measurement can easily be obtained in clinical routine. We consider that the measurements are noisy, and use the same observation model as Mosca et al. [15]. For mature cells, it is assumed a Gaussian noise with a variance that depends on the true VAF $y_k^{(i)} \in [0, 1]$ at time $t_k^{(i)}$:

$$\hat{y}_k^{(i)} | y_k^{(i)} \sim \mathcal{N}\left(y_k^{(i)}, y_k^{(i)}(1 - y_k^{(i)})\sigma_m^2\right) \quad (\text{A.1})$$

with σ_m to be estimated.

The clonal architecture of purified CD34⁺ progenitors is measured at different time points. To determine this clonal architecture from a blood sample, several experimental steps are required, as presented in [15]. Such measurements cannot be easily obtained in clinical routine, especially not every three or four months. When measuring the clonal architecture at time $t_k^{(i)}$, we get $\hat{n}_{k,wt}^{(i)}$, $\hat{n}_{k,hom}^{(i)}$, and $\hat{n}_{k,hom}^{(i)}$ of wt (wild-type), heterozygous (het) and homozygous (hom) mutated progenitor cells, respectively. From these values, we can compute the mutated heterozygous Clonal Fraction (CF):

$$\hat{z}_{k,hom}^{(i)} = \frac{\hat{n}_{k,hom}^{(i)}}{\hat{n}_{k,wt}^{(i)} + \hat{n}_{k,hom}^{(i)} + \hat{n}_{k,hom}^{(i)}}$$

as well as the mutated homozygous CF:

$$\hat{z}_{k,hom}^{(i)} = \frac{\hat{n}_{k,hom}^{(i)}}{\hat{n}_{k,wt}^{(i)} + \hat{n}_{k,hom}^{(i)} + \hat{n}_{k,hom}^{(i)}}$$

The uncertainty associated with the measured CF comes from the fact that only a limited number of progenitor cells are sampled. Following what was done in Mosca et al. [15] (and also in [1]), we consider that the observed immature cells are randomly sampled from an unknown but very large set of immature cells so that the uncertainty might be modeled by a multinomial distribution:

$$\mathbb{P}\left[\hat{n}_{k,hom}^{(i)} = n_1, \hat{n}_{k,hom}^{(i)} = n_2, \hat{n}_{k,hom}^{(i)} = n_3 \mid z_{k,hom}, z_{k,hom}\right] = \frac{(n_1 + n_2 + n_3)!}{n_1!n_2!n_3!} z_{k,wt}^{n_1} z_{k,hom}^{n_2} z_{k,hom}^{n_3}$$

where $z_{k,wt}$, $z_{k,hom}$, and $z_{k,hom}$ are the true CF for wt, het, and hom progenitor cells, respectively (with $z_{k,wt} = 1 - z_{k,hom} - z_{k,hom}$).

Both the CF and the VAF are ratios whose values range between 0 and 1. The CF refers to a ratio of mutated cells, while the VAF refers to a ratio of mutated alleles. Implicitly, the VAF and the CF refer to mature and progenitor cells, respectively. We will explicitly specify when the VAF and the CF refer to the HSCs.

Finally, $\mathcal{I}^{(i)}$ corresponds to the set of observation times for patient i , and we have:

$$\mathcal{D}^{(i)} = \left(t_k^{(i)}, \hat{n}_{k,wt}^{(i)}, \hat{n}_{k,hom}^{(i)}, \hat{n}_{k,hom}^{(i)}, \hat{y}_k^{(i)}\right)_{k \in \mathcal{I}^{(i)}}$$

Clinical and experimental data for each of the 19 $JAK2^{V617F}$ MPN patients are represented in Figures A.1 to A.19. The patients' ids are those used by Mosca et al. [15]. In this article, we will focus on the three patients presented in Tab. A.1. These are the patients that will be excluded from the cohort of the 19 patients, each in turn, such that we could study for them:

- our capacity to predict their long-term responses to IFN α ,

	#12	#18	#32
Age at t_0 [years]	63	57	52
Disease	PV	PMF	PV
Sex	F	F	M

Table A.1: List of the patients we focus on in this study. For each of them, we indicate their age at the beginning of the therapy (t_0), their disease (PV: polycythemia vera; PMF: primary myelofibrosis), and whether they are a female (F) or a male (M). The patients' numbers are the same as in the article of Mosca et al. [15]. More information about the patients can be found in their study.

- how to optimize their therapy,
- and what would have been the optimal experimental design.

Suppose we consider, for example, patient #32. When we focus on this patient, we exclude them from the cohort. Thus, we also do not consider all of their CF measurements, but only two of them: one at the diagnostic time and one later. We exclude the other CF observations to mimic the fact that this patient would be observed following a realistic clinical routine. The excluded data will be used to assess the quality of the predictions. However, when we focus on another patient, for example, patient #12, then patient #32 is part of the cohort. In this case, we consider all the observations for patient #32.




Figure A.1: Longitudinal observations of patient #2. The x-axis corresponds to the time from the start of the IFN α therapy ($t = 0$). Clinical data consist of the measurements of the VAF in mature cells (black squares), Experimental data to the heterozygous (green circles) and homozygous (blue triangles) CF in progenitor cells. The shaded beige areas correspond to the weekly dose of IFN α received over time. Some observations might have been obtained before the start of the therapy.




Figure A.2: Longitudinal observations of patient #3.




Figure A.3: Longitudinal observations of patient #6.




Figure A.4: Longitudinal observations of patient #8.




Figure A.5: Longitudinal observations of patient #11.




Figure A.6: Longitudinal observations of patient #12.




Figure A.7: Longitudinal observations of patient #15.

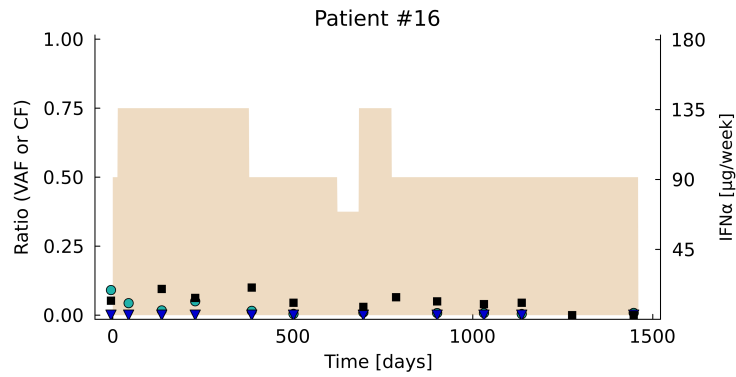


Figure A.8: Longitudinal observations of patient #16.




Figure A.9: Longitudinal observations of patient #18.




Figure A.10: Longitudinal observations of patient #19.




Figure A.11: Longitudinal observations of patient #20.




Figure A.12: Longitudinal observations of patient #22.




Figure A.13: Longitudinal observations of patient #23.




Figure A.14: Longitudinal observations of patient #25.




Figure A.15: Longitudinal observations of patient #27.




Figure A.16: Longitudinal observations of patient #29.




Figure A.17: Longitudinal observations of patient #30.




Figure A.18: Longitudinal observations of patient #31.




Figure A.19: Longitudinal observations of patient #32.

B Model

The model on which we base this study was first proposed by Mosca et al. [15] to describe the dynamics of mutated cells over IFN α therapy and then extended in [9] to account for dose variations during the treatment. In the original formulation from Mosca et al., three populations of cells were considered according to whether the cells are wild-type (WT) or have the $JAK2^{V617F}$ mutation in one (heterozygous, subscript *het*) or two (homozygous, subscript *hom*) alleles. These three populations were assumed to be independent. If we consider one of them (Fig. 2 - main article), their cell dynamics were described by the following system of ordinary differential equations (for clarity, we omit the subscript corresponding to the population):

$$\begin{cases} \dot{N}_q(t) &= -\gamma N_q(t) + \beta N_a(t) \\ \dot{N}_a(t) &= \gamma N_q(t) + (\alpha\Delta - \beta)N_a(t) \\ \dot{N}_i(t) &= \alpha(1 - \Delta)\kappa_i N_a(t) - \delta_i N_i(t) \\ \dot{N}_m(t) &= \delta_i \kappa_m N_i(t) - \delta_m N_m(t) \end{cases} \quad (\text{B.1})$$

where $N_q(t)$, $N_a(t)$, $N_i(t)$, and $N_m(t)$ describe respectively the numbers of quiescent HSCs, active HSCs, progenitors (that we also call immature cells) and mature cells. $\gamma N_q(t)$ and $\beta N_a(t)$ correspond to the rate at which HSCs become respectively active or quiescent, α is the division rate of active HSCs, Δ corresponds to the balance between differentiated and symmetrical division, κ_i and κ_m are the proliferative rates of immature and mature cells, respectively, δ_i is the rate constant controlling how progenitors exit their compartment to become mature cells with a death rate δ_m .

IFN α was found to act on the differentiation and quiescence exit of mutated stem cells, that is, on parameters γ and Δ [15]. In a previous study [9], after conducting a model selection procedure, we found that an affine sigmoid relation can model the impact of the variations of IFN α dose $d(t)$ on parameters Δ_{het} and Δ_{hom} through:

$$\Delta_{het} : d \mapsto -2 \left(\frac{1}{1 + e^{-\rho_{het} \cdot d}} - 0.5 \right) \cdot (1 + \delta_{0,het}) + \delta_{0,het} \quad (\text{B.2})$$

and equivalently for homozygous cells. The dose $d(t) \in [0, 1]$ is defined as the normalized weekly averaged IFN α quantity administrated to the patient. For normalization, we divide by 180 μ g/week, considered the maximally tolerated dose - this maximum value seems to be a consensus but not the result of any clinical trial. $d(t)$ is a piece-wise constant input function in our model. After several biological assumptions detailed in [15, 9], the parameters to be estimated for each patient are the following:

- η_{het} and η_{hom} : the proportions of heterozygous and homozygous mutated HSCs, respectively, at the beginning of the treatment,
- γ_{het} and γ_{hom} : the rate constants of HSC quiescence exit during IFN α therapy,
- $\delta_{0,het}$, $\delta_{0,hom}$, ρ_{het} and ρ_{hom} , which describe how parameters Δ_{het} and Δ_{hom} behave depending on the IFN α dose (see eq. (B.2)). Δ_{het} and Δ_{hom} model to which extent mutated HSCs differentiate, therefore leading to an exhaustion of the mutated stem cell pool,
- $\kappa_{m,het} = \kappa_{m,hom}$, which models the proliferative advantage of mutated cells at the last stages of hematopoiesis.

C Parameter estimation

We have the data of $N = 19$ *JAK2*^{V617F} MPN patients, as presented in section A. We denote by \mathcal{D} the dataset for all 19 patients. Let us consider a patient i that we will focus on. We denote by $\mathcal{D}^{(i)}$ their dataset, and by $\mathcal{D}_{-i} = \mathcal{D} \setminus \mathcal{D}^{(i)}$ the dataset of all patients except individual i .

C.1 Hierarchical inference

The hierarchical Bayesian inference method is detailed in our previous study [9]. We provide here the information necessary for a good understanding of the methods.

$\boldsymbol{\theta}^{(-i)} = \{\boldsymbol{\theta}^{(j)}\}_{1 \leq j \leq N, j \neq i}$ denotes the set of all patient parameters - except patient i - with:

$$\boldsymbol{\theta}^{(j)} = \left(\theta_1^{(j)}, \dots, \theta_P^{(j)} \right)$$

where $P = 9$ is the number of parameters to estimate. We assume that all (except η_{het} and η_{hom}) individual parameter vectors are realizations of the same random variable of a distribution of unknown mean and variance:

$$\forall 1 \leq j \leq N, j \neq i, \forall k \in \{1, \dots, P\}, \theta_k^{(j)} | \tau_k^{(-i)}, \sigma_k^2^{(-i)} \sim \mathcal{N}_{c,k} \left(\tau_k^{(-i)}, \sigma_k^2^{(-i)} \right) \quad (\text{C.1})$$

where the population distribution for each component is a truncated Gaussian distribution $\mathcal{N}_{c,k}$ over a range that depends on the considered parameter k (see next paragraph), and $\boldsymbol{\tau}^{(-i)} = (\tau_1^{(-i)}, \dots, \tau_P^{(-i)})$ and $\boldsymbol{\sigma}^2^{(-i)} = (\sigma_1^2^{(-i)}, \dots, \sigma_P^2^{(-i)})$ are the hyper-parameters. The prior distributions of the hyper-parameters $\tau_k^{(-i)}$ are chosen uniform, over the same ranges as for the parameters they are associated with. For each k , $\sigma_k^2^{(-i)}$ follows an improper prior distribution - namely, an inverse-gamma (0,0) law. Then, we can estimate the joint posterior distributions of $\boldsymbol{\theta}^{(-i)}$ and hyper-parameters $\boldsymbol{\tau}^{(-i)}$ and $\boldsymbol{\sigma}^2^{(-i)}$, as done in [15] and detailed in [9]:

$$\begin{aligned} p \left[\boldsymbol{\theta}^{(-i)}, \boldsymbol{\tau}^{(-i)}, \boldsymbol{\sigma}^2^{(-i)} \mid \mathcal{D}_{-i} \right] &\propto p \left[\mathcal{D}_{-i} \mid \boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}^{(i+1)}, \dots, \boldsymbol{\theta}^{(N)}, \boldsymbol{\tau}^{(-i)}, \boldsymbol{\sigma}^2^{(-i)} \right] \times \\ & p \left[\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}^{(i+1)}, \dots, \boldsymbol{\theta}^{(N)}, \boldsymbol{\tau}^{(-i)}, \boldsymbol{\sigma}^2^{(-i)} \right] \\ &\propto \prod_{1 \leq j \leq N, j \neq i} \left(p \left[\mathcal{D}^{(j)} \mid \boldsymbol{\theta}^{(j)} \right] p \left[\boldsymbol{\theta}^{(j)} \mid \boldsymbol{\tau}^{(-i)}, \boldsymbol{\sigma}^2^{(-i)} \right] \right) p \left[\boldsymbol{\tau}^{(-i)} \right] p \left[\boldsymbol{\sigma}^2^{(-i)} \right] \end{aligned}$$

Then, we sample from the posterior distribution using a Markov Chain Monte Carlo (MCMC) method, namely the Metropolis-Hastings within Gibbs algorithm [7, 3]. Conditionally on the hyper-parameters, patients are independent, and their parameters can be sampled using a standard Metropolis-Hasting scheme. The hyper-parameters are sampled using the Gibbs method, which consists of sampling from the marginal conditional posterior distribution of the hyper-parameters. Details of the calculations are presented in [9].

To initialize the MCMC chain, we first run an optimization algorithm - namely the CMA-ES [4, 5]. The CMA-ES is a stochastic algorithm that we use for looking at the estimator of the maximum *a posteriori*. We then use this estimator for initializing the MCMC. Furthermore, the CMA-ES algorithm also learns a covariance matrix that we choose (up to a multiplication factor) for our proposal in the Metropolis-Hasting scheme. Model calibration was achieved by implementing the previous methods in the Julia programming language. The framework used for parameter estimation is available at:

<https://gitlab-research.centralesupelec.fr/2012hermangeg/bayesian-inference>

We run our hierarchical inference method over 13 million iterations, with a burn-in length of 2 million, until achieving convergence.

After running the estimation procedure, we have for $\tau_k^{(-i)}$ and $\sigma_k^2^{(-i)}$ the estimate of their posterior distribution from which we estimate the posterior mean, given the observations:

$$\begin{aligned} \bar{\tau}_k^{(-i)} &= \mathbb{E}[\tau_k^{(-i)} \mid \mathcal{D}_{-i}] \\ \bar{\sigma}_k^2^{(-i)} &= \mathbb{E}[\sigma_k^2^{(-i)} \mid \mathcal{D}_{-i}] \end{aligned}$$

These latter values will be used as prior information when considering a new patient.

C.2 Parameter estimation for a new patient

Patient i has been excluded from the cohort of patients. This patient can therefore be considered a new patient. As presented in section A, for them, we consider having only two observations of the clonal architecture, similar to what could be obtained in clinical routines.

C.2.1 Prior

$\boldsymbol{\theta}^{(i)}$ is the parameter vector of patient i :

$$\boldsymbol{\theta}^{(i)} = \left(\theta_1^{(i)}, \dots, \theta_P^{(i)} \right)$$

For η_{het} (and η_{hom}) which correspond to the initial quantities of heterozygous (and homozygous) mutated HSCs compared to the number of wild-type HSCs, we consider a prior uniform over $[0, 3]$ (unitless parameter). For the other parameters (index k), we consider *a priori*:

$$\theta_k^{(i)} \sim \mathcal{N}_{c,k} \left(\bar{\tau}_k^{(-i)}, \bar{\sigma}_k^2{}^{(-i)} \right) \quad (\text{C.2})$$

$\mathcal{N}_{c,k}$ is a truncated Gaussian distribution over a range that depends on the considered parameter :

- For parameters $\delta_{0,het}$ and $\delta_{0,hom}$, the range is $[0.0, 0.5]$ (unitless parameters),
- For ρ_{het} and ρ_{hom} , the range is $[0, 10]$ (unitless parameters),
- For k_m , the range is $[1, 20]$ (unitless parameter),
- For $1/\gamma_{het}$ and $1/\gamma_{hom}$, the range is $[10, 300]$ (in days).

C.2.2 Bayesian Inference

We denote by $\mathcal{I}_T^{(i)}$ the set of observation times before the assimilation time T . The dataset of patient i used for estimating their parameters is:

$$\mathcal{D}_T^{(i)} = \left\{ t_k^{(i)}, \hat{y}_k^{(i)} \right\}_{k \in \mathcal{I}_T^{(i)}} \cup \left\{ \hat{n}_{k=1,wt}^{(i)}, \hat{n}_{k=1,het}^{(i)}, \hat{n}_{k=1,hom}^{(i)} \right\} \cup \left\{ \hat{n}_{k',wt}^{(i)}, \hat{n}_{k',het}^{(i)}, \hat{n}_{k',hom}^{(i)} \right\}$$

with k' corresponding to the index of the observation made at a time close to 300 days after treatment. The posterior distribution is expressed by:

$$p[\boldsymbol{\theta}_T^{(i)} | \mathcal{D}_T^{(i)}] \propto p[\mathcal{D}_T^{(i)} | \boldsymbol{\theta}_T^{(i)}] p[\boldsymbol{\theta}_T^{(i)}] \quad (\text{C.3})$$

with $p[\boldsymbol{\theta}_T^{(i)}]$ the prior expressed in the previous paragraph and $p[\mathcal{D}_T^{(i)} | \boldsymbol{\theta}_T^{(i)}]$ the likelihood.

First, we find the parameter vector $\hat{\boldsymbol{\theta}}_T^{(i)}$ which maximizes the previous posterior. $\hat{\boldsymbol{\theta}}_T^{(i)}$ is estimated using the CMA-ES algorithm [4]. Then, we run the Metropolis-Hasting algorithm, whose proposal is set to the covariance matrix learnt using the CMA-ES, and the initial state set to $\hat{\boldsymbol{\theta}}_T^{(i)}$. We run the Bayesian inference method over 1 million iterations, and verify that we achieved convergence.

In the prediction part of this study, we will study the posterior distribution of $\boldsymbol{\theta}_T^{(i)}$ according to the value of the assimilation time T .

When studying the treatment optimization or the experimental design, we consider the assimilation time $T = 600$ days, and we will set the parameter vector to the mean posterior value $\mathbb{E}[\boldsymbol{\theta}_{T=600}^{(i)} | \mathcal{D}_{T=600}^{(i)}]$.

C.2.3 Uncertainty propagation

In the prediction part of this study, we infer not only the parameter posterior distribution but also the dynamics of the mutated cells (or, more precisely, the evolution of the VAF and the CF over time). For this purpose, we sample from the posterior distribution some parameter vectors using a Monte-Carlo method and propagate the uncertainty from the input parameters to the output of the model (that is, the dynamics of the CF in each hematopoietic compartment).

C.2.4 Estimating d_{inf}

In our model, the parameter Δ_{het} depends on the dose $d \in [0, 1]$ according to the equation (B.2):

$$\Delta_{het} : d \mapsto -2 \left(\frac{1}{1 + e^{-\rho_{het} \cdot d}} - 0.5 \right) \cdot (1 + \delta_{0,het}) + \delta_{0,het}$$

The same goes for Δ_{hom} .

The parameter $\Delta_{het} \in [-1, 1]$ corresponds to the balance between the differentiated and symmetrical divisions of mutated HSCs. $\Delta_{het} > 0$ means that the mutated heterozygous clone continues to expand, that is, to invade the stem cell pool. $\Delta_{het} < 0$ means that mutated HSCs encounter more differentiated divisions than symmetrical ones, so the pool of mutated HSCs will be depleted.

The necessary condition for the pool of mutated heterozygous HSCs to be potentially totally eradicated in the future is to use a dose d such that $\Delta_{het}(d) < 0$. According to our model, Δ_{het} decreases with the dose. Thus, there exists $d_{inf,het}$ such that:

$$\Delta_{het}(d_{inf,het}) = 0$$

$$\Delta_{het}(d) > 0 \text{ for } d < d_{inf,het}$$

The same goes for homozygous cells. We define $d_{inf} = \max(d_{inf,het}, d_{inf,hom})$ as the minimal dose to administer to the patient such that both the homozygous and heterozygous mutated cell could be depleted.

d_{inf} depends on parameters $\delta_{0,het}$, $\delta_{0,hom}$, ρ_{het} , and ρ_{hom} . Therefore, d_{inf} is patient-dependent.

C.2.5 Assessing the quality of the fits and predictions

For a given assimilation time T , we have estimated the model parameters (that is, their posterior distribution) based on the observations included in the dataset $\mathcal{D}_T^{(i)}$. After having propagated the uncertainties, we get the dynamics of the VAF among mature cells and the CF among progenitors. That is, for a given time, we get a distribution of the values of the VAF, and the heterozygous and homozygous CF, from which we compute (and display) the median. For each clinical and experimental data (either a CF or a VAF measurement), observed at a time t , we can then confront the experimental value to the theoretical one (i.e., the median value introduced above). For a given observation time, the distance between the experimental and the theoretical value will be the quadratic error. This error can be computed either for a VAF measurement, or an heterozygous CF, or an homozygous CF. When considering several observations, the reported error will then be the mean square error (MSE). First, the MSE could be computed based on the observations from $\mathcal{D}_T^{(i)}$. In that case, we assess how well the model fits the data (used for its calibration). The MSE can also be computed based on observations coming from the control dataset: $\mathcal{D}_c^{(i)} = \mathcal{D}^{(i)} \setminus \mathcal{D}_T^{(i)}$ (which includes all observations not used for the model calibration). In that case, we assess how well the models gives accurate predictions.

D Interrupting the therapy

This section considers the case when no IFN α is administered to the patient. It is the case when the treatment has not yet started or has been interrupted.

D.1 Criterion

We aim at optimizing the therapy from a time T to a time τ ; this latter time corresponds to the end of the treatment. IFN α therapy is a long-term treatment that clinicians do not want to prolong indefinitely due to adverse effects in most cases. We aim to find a criterion for deciding when the therapy could be stopped. Such a criterion should be acceptable both for clinicians and patients.

Since MPNs are diseases that develop over a long time [17, 16], it is reasonable to choose to stop treatment even before complete eradication of the mutated HSCs, if there is a good reason to believe that the clonal development following the discontinuation of treatment is unlikely to lead to the disease recurrence (that is, the reappearance of the symptoms) during the patient's lifetime. In this work, we propose to determine, according to different criteria, when one could suggest to the clinician (and the patient) to stop the treatment.

Unless complete remission is achieved (which is difficult, if not impossible, to estimate clinically), the patient may still have several mutated stem cells prone to revert to clonal development leading to a recurrence of symptoms. Nevertheless, because of the slow clonal expansion, it can be considered that if the proportion of mutated cells remains below the threshold for symptom onset during the patient's lifetime, this will not be a problem. Of course, one cannot know how long the patient will live. One can make a conservative choice by choosing an age limit beyond 100 years, for example, which would mean wanting an almost zero risk of a possible relapse of the patient. However, such a choice might not be relevant, given the life expectancy of individuals. We thus choose to set the age limit (below which we do not want the threshold to be reached) at 90 years. This age corresponds to the life expectancy of 65-year-old individuals in 2030 in developed countries [11].

Concerning the threshold of onset of the disease, we had assumed in a previous study [8] that the disease could appear for a CF of heterozygous mutated HSCs (in patients having no homozygous clones) of 15%. That is, a VAF of 7.5% in HSCs. This latter threshold is based on the study by Dupont et al. [2].

Considering a life expectancy, at 65 years old, equal to 90 [11], we will consider interrupting the treatment of a patient at an age τ such that $\phi_\tau(t = 90 \text{ years}) = 7.5\%$ with:

$$\begin{aligned} \phi_\tau(t) &= \frac{0.5\tilde{N}_{het}(t) + \tilde{N}_{hom}(t)}{\tilde{N}_{het}(t) + \tilde{N}_{hom}(t) + N_{WT}} \\ &= \frac{0.5\tilde{N}_{het}(\tau)e^{s_{het}(t-\tau)} + \tilde{N}_{hom}(\tau)e^{s_{hom}(t-\tau)}}{\tilde{N}_{het}(\tau)e^{s_{het}(t-\tau)} + \tilde{N}_{hom}(\tau)e^{s_{hom}(t-\tau)} + N_{WT}} \end{aligned}$$

which corresponds to the VAF, among HSCs, at time t , given the quantities (or the equivalent CF) of mutated HSCs at time τ . N_{WT} is the number of wild-type HSCs, assumed to be constant and equal to 10^5 [12]. More details about the quantities involved in the expression of $\phi_\tau(t)$ are presented in the next paragraphs.

To sum up, we will choose to stop treatment at an age such that the subsequent clonal development will not lead to a VAF of 7.5% (among HSCs) before the age of 90.

D.2 Model of the clonal expansion without IFN α

First, we consider a patient having only heterozygous $JAK2^{V617F}$ clones. We assume that the clonal development after interruption of the treatment follows the same dynamics as before the treatment and that this dynamic is described by the model proposed in [8]. We assume that we remain in the case where the number of mutated cells remains sufficient for the following deterministic approximation - derived in [8] - to be valid:

$$\tilde{N}_{het}(t) = \tilde{N}_{0,hct} \exp(s_{het}(t - t_0)) \quad (\text{D.1})$$

where the fitness of the heterozygous clone was estimated to be equal to $s_{het} = 20.4\%/year$. Here, $\tilde{N}_{het}(t) = \tilde{N}_{a,het}(t) + \tilde{N}_{q,het}(t)$ corresponds to the total number of heterozygous HSCs (active and quiescent). Equation (D.1) is valid during the MPN development in the absence of IFN α therapy. We use the tilde symbol $\tilde{\circ}$ to denote a variable which dynamics is studied when the patient is not under treatment (that is, either before the start of the treatment, or after the patient has definitively interrupted the therapy).

A sufficient condition for the validity of the deterministic approximation was shown to be $\tilde{N}_{0,het} = 2,000$, but this latter choice was conservative in the sense that, even for lower values of $\tilde{N}_{0,het}$, the approximation error was low [8].

Now, we consider the more general case of a patient having both heterozygous and homozygous clones that can develop in parallel. In that case, the heterozygous CF among HSCs is defined by:

$$\begin{aligned} CF_{het}(t) &= \frac{\tilde{N}_{het}(t)}{\tilde{N}_{het}(t) + \tilde{N}_{hom}(t) + N_{WT}} \\ &= \frac{\tilde{N}_{q,het}(t) + \tilde{N}_{a,het}(t)}{\tilde{N}_{q,het}(t) + \tilde{N}_{a,het}(t) + \tilde{N}_{q,hom}(t) + \tilde{N}_{a,hom}(t) + N_{WT}} \end{aligned} \quad (D.2)$$

And the VAF among HSCs:

$$VAF(t) = \frac{0.5 \tilde{N}_{het}(t) + \tilde{N}_{hom}(t)}{\tilde{N}_{het}(t) + \tilde{N}_{hom}(t) + N_{WT}} \quad (D.3)$$

In all this study, we consider that the number of the WT HSCs is constant and equal to $N_{WT} = 10^5$ [12]. In particular, we do not assume that, during the clonal expansion, mutated cells might replace WT cells. On the contrary, we consider that mutated HSCs expand in parallel to the WT, potentially conquering other hematopoietic niches [6], such that the number of WT (normal) HSCs remains constant. To note that the hypothesis of a constant number of WT HSCs during a clonal expansion of mutated HSCs was also done by Michor et al. [14] when studying the dynamics of chronic myeloid leukemia. We assume that the expansion of the malignant homozygous mutated cells will follow the same dynamics as the heterozygous one:

$$\tilde{N}_{hom}(t) = \tilde{N}_{0,hom} \exp(s_{hom}(t - t_0)) \quad (D.4)$$

with a fitness s_{hom} higher than for the heterozygous case [17]. We estimate the value of the homozygous fitness in the next section.

Then, we get the expression of the VAF among HSC over time $t \geq t_0$, given the values at time t_0 :

$$VAF(t) = \frac{0.5 \tilde{N}_{het}(t_0)e^{s_{het}(t-t_0)} + \tilde{N}_{hom}(t_0)e^{s_{hom}(t-t_0)}}{\tilde{N}_{het}(t_0)e^{s_{het}(t-t_0)} + \tilde{N}_{hom}(t_0)e^{s_{hom}(t-t_0)} + N_{WT}} \quad (D.5)$$

D.3 Estimating the fitness of the $JAK2^{V617F}$ homozygous clone

D.3.1 Approach

To describe the clonal expansion without IFN α , we use the model we previously studied in [8]. From our previous work, we also estimated the fitness of the heterozygous clone to be equal to $s_{het} = 20.4\%/year$. However, the study was not conducted for patients with homozygous clones, so we do not have an estimation of s_{hom} . We expect to have $s_{hom} > s_{het}$ [17]. We consider that:

$$s_{hom} = \beta s_{het}$$

with $\beta > 1$. To estimate s_{hom} (or, equivalently, β), we will use the data from Williams et al. [17] (Fig. D.1). In their study, they had some patients with heterozygous mutated clones and some others with homozygous mutated clones. Using a model they calibrated based on the whole genome sequencing of progenitor cells and the subsequent construction of phylogenetic trees, they inferred the fitness S_i of mutated $JAK2^{V617F}$ clones for different patients (index i). From their study, we can get their estimated $S_{i,het}$ for $1 \leq i \leq n_{het}$ and $S_{i,hom}$ for $1 \leq i \leq n_{hom}$, with n_{het} and n_{hom} the number of heterozygous and homozygous patients they studied, respectively. If we note \tilde{S}_{het} the median

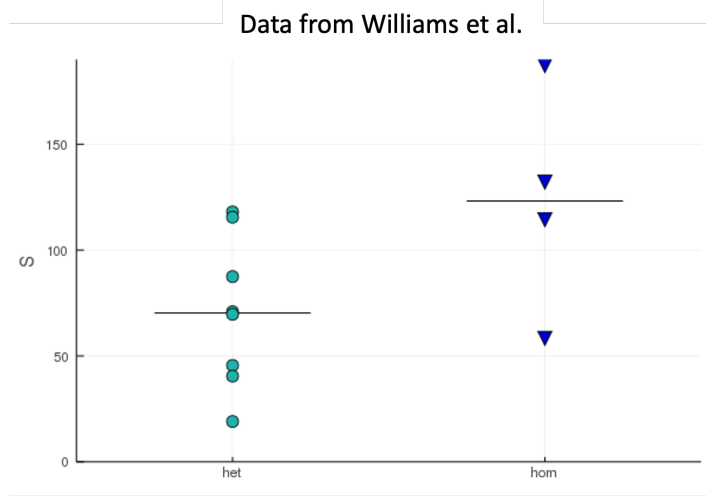


Figure D.1: Data from Williams et al. [17]. From their article, we extracted the mean value of their so-called S parameter, which corresponds to the fitness of a mutated clone. We display the value of S according to whether it refers to a $JAK2^{V617F}$ heterozygous clone or a homozygous one. There is a trend to have higher fitness in the homozygous condition ($p = 0.1535$ with a Mann-Whitney test [13]).

heterozygous fitness and \tilde{S}_{hom} the median homozygous fitness, we could compute $\beta' = \tilde{S}_{hom}/\tilde{S}_{het}$, and, in a naive approach, set $\beta = \beta'$. Yet, this first estimation of β would be too approximate since their model differs significantly from ours. Indeed, Williams et al. assume the HSC pool to be of constant size, and, therefore, that mutated cells take the place of WT cells during their expansion. On the contrary, we assume in our model that we have a constant number of WT cells and that the mutated cells expand without replacing the former. To circumvent the difficulty related to the fact that our models are different, we will try to find a relation $f : S_{het,i} \mapsto s_{het,i}$ such that $s_{het,i}$ is consistent with the fitness used in our model. Then, we could find a better estimation for β using:

$$\beta \approx \frac{f(\tilde{S}_{het})}{f(\tilde{S}_{hom})} \quad (\text{D.6})$$

D.3.2 Finding the relation $f : S_{het} \mapsto s_{het}$

Here, we will approximate a relation that links the so-called "fitness of clone" S of Williams et al. [17] to our fitness s . In this paragraph, we consider the case of heterozygous mutated cells and omit the subscript *het* for clarity. In Williams et al.'s model, the number of mutated cells $N(t)$ follows the differential equation:

$$\frac{dN}{dt} = N(t)\alpha\varsigma \left(1 - \frac{N(t)}{N_{tot}}\right) \quad (\text{D.7})$$

and they define:

$$S = \exp(\alpha\varsigma) - 1$$

where ς is the selective advantage (in their paper, they note it actually s but, to avoid the confusion with our fitness parameter, we note it ς), α corresponds to a division rate as in [8], and N_{tot} is the total size of the HSC pool (WT and mutated) assumed to be constant. In our case, on the contrary, we have $N_{tot} = N_{WT} + N(t)$ with $N(t)$ the number of mutated cells. In our model, if we differentiate equation (D.1), we get:

$$\frac{dN}{dt} = s N(t) \quad (\text{D.8})$$

By equalizing equations (D.7) and (D.8), we get:

$$s = \log(1 + S) \left(1 - \frac{N(t)}{N_{tot}}\right)$$

If we assume that the CF among HSCs $N(t)/N_{tot}$ is approximately equal to the one among progenitors η (that is, the CF that can be measured by Williams et al.), we get:

$$s \approx \log(1 + S)(1 - \eta(t)) \quad (\text{D.9})$$

If we integrate the previous equation from the acquisition time to the sampling time, and assume a linear clonal expansion (which we acknowledge is a rough approximation), we get:

$$s \approx \log(1 + S) \left(1 - \frac{\hat{\eta}}{2}\right) = f(S) \quad (\text{D.10})$$

with $\hat{\eta}$ the CF among progenitor cells measured at the observation time.

D.3.3 Results

For each individual i from the cohort of Williams et al. [17], we have their estimate S_i and we can estimate the progenitor CF $\hat{\eta}_i$ at sampling time by counting the numbers of mutated and WT individual colonies at the tips of the phylogenetic tree branches. Below, we report the values from Williams et al. that we use for this analysis. The data for the patients having heterozygous but no (or negligible) homozygous clones are the following:

- PD6646: $S = 119$ %/year (clade JAK2, DNMT3A) - $\hat{\eta} = 56.7\%$
- PD5179: $S = 116$ %/year (clade JAK2, 1q+, 9q-, 9+) - $\hat{\eta} = 96.7\%$
- PD6629: $S \approx 66$ %/year (clade DNMT3A, JAK2, TET2, and clade DNMT3A, JAK2) - $\hat{\eta} = 43.8\%$
- PD7271: $S = 73$ %/year (clade JAK2) - $\hat{\eta} = 21.6\%$
- PD9478: $S = 71$ %/year (clade JAK2, DNMT3A) - $\hat{\eta} = 89.5\%$
- PD5163: $S = 45$ %/year (clade JAK2) - $\hat{\eta} = 8.6\%$
- PD5117: $S = 18$ %/year (clade JAK2) - $\hat{\eta} = 55.7\%$

For patient PD6629, Williams et al. have two estimations of S according to the considered clade ($S = 90$ %/year for the clade DNMT3A, JAK2, TET2, and $S = 41$ for the clade DNMT3A, JAK2), we consider the mean value.

For patients having homozygous clones but no (or negligible) heterozygous ones, the data are the following:

- PD4781: $S=114\%$ /year (clade TET2, JAK2, 9pUPD, 7p-, 7p+) - $\hat{\eta} = 97.44\%$
- PD5847: $S \approx 122.5$ %/year (clade TET2, JAK2, 9pUPD and clade JAK2, 9pUPD) - $\hat{\eta} = 80.2\%$
- PD5182: $S = 132$ %/year (clade 9pUPD, JAK2) - $\hat{\eta} = 34.38\%$

For patient PD5847, Williams et al. have two estimations of S according to the considered clade ($S = 187$ %/year for the clade TET2, JAK2, 9pUPD, and $S = 58$ for the clade JAK2, 9pUPD), we consider the mean value.

From the previous values and using eq. (D.10), we can compute a fitness $s_i = f(S_i)$ for each patient i from Williams et al. [17]. Results are displayed in figure D.2.

Then, we estimate our parameter β by computing the ratio between the median of the s values for homozygous patients ($\tilde{s}_{hom} = 47.9$ %/year), and the median of the s values for the heterozygous patients ($\tilde{s}_{het} = 39.6$ %/year). We find $\beta = 1.21$.

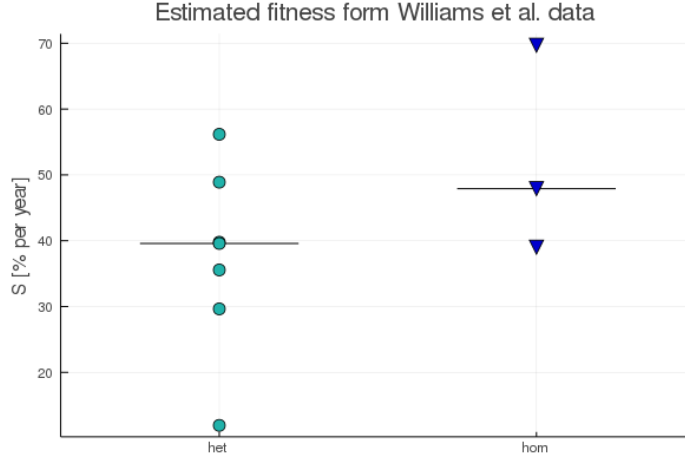


Figure D.2: Estimation of a fitness value s for each patient from Williams et al. [17] using eq. (D.10).

E Optimizing the therapy

E.1 Grid search of the optimal parameters

For a given patient, we estimated their parameter vector θ as presented in § C.2.2, here with the assimilation time $T = 600$ days. The parameter vector entirely describes the patient’s response to the treatment, for any potential dose $d(t) \in [0, 1]$ over time t . We consider three therapeutic strategies (index s), as presented in §2.3.2 in the main text:

- $s = 1$ corresponding to the constant strategy,
- $s = 2$ corresponding to the periodic strategy,
- $s = 3$ corresponding to the decreasing strategy (see also § E.3).

We note $\mathbf{X} \in \mathbb{R}^3$ the parameter vector of dimension 3 associated with the strategy:

$$\mathbf{X} = (\bar{d}, L, \lambda)$$

As we have defined it - and without any constraints on parameters L and λ - the constant strategy would be a special case of both the decreasing and periodic strategies. The constant strategy is the limiting case of the decreasing strategy when $\lambda = 1$, and the limiting case of the periodic strategy for $L \rightarrow +\infty$. In practice, given that we study times on the scale of human life, $+\infty$ could be replaced by any high enough time value so that we would still have $\mathbf{X} \in \mathbb{R}^3$.

The choice of \mathbf{X} entirely characterizes which dose would be administered to the patient for $t \geq T$, if the treatment were not interrupted. That is:

$$d_s : \begin{cases} \mathbb{R}^3 \times [T, +\infty[& \longrightarrow & [0, 1] \\ (\mathbf{X}, t) & \longmapsto & d_s(t) \end{cases}$$

Then, the values of θ and \mathbf{X} entirely characterize the response to the treatment for any time $t \geq T$. In particular, we infer the dynamics of the VAF among HSCs. By applying the criterion presented in § D, we can estimate the time τ when the treatment could be interrupted. To note that, to apply this criterion, we also need to know the age of the patient at the beginning of their therapy (see Tab. A.1). The interruption time is thus a function:

$$\tau_s : \begin{cases} \mathbb{R}^3 & \longrightarrow & [T, +\infty[\\ \mathbf{X} & \longmapsto & \tau \end{cases}$$

We omit to precise that τ also depends on the patient’s parameter vector (and their age) since these latter values are supposed to be known, and because they do not come into play in the optimization problem we study here.

We consider four scenarios of dose-toxicity (index x) as presented in §2.3.3 in the main text:

- $x = 1$ corresponding to the linear relation,
- $x = 2$ corresponding to the convex relation,
- $x = 3$ corresponding to the concave relation,
- $x = 4$ corresponding to the composite relation (see also § E.4).

We define the toxicity as a function of the dose:

$$z_x : \begin{cases} [0, 1] & \longrightarrow \mathbb{R}^+ \\ d & \longmapsto z_x(d) \end{cases}$$

Finally, we define the toxicity-related amount of IFN α administered from T to the end of the therapy:

$$M_x : \begin{cases} \mathbb{R}^3 & \longrightarrow \mathbb{R}^+ \\ \mathbf{X} & \longmapsto \int_T^{\tau_s(\mathbf{X})} z_x(d_s(\mathbf{X}, t)) dt \end{cases} \quad (\text{E.1})$$

Our optimization problem to solve is to find the parameter vector \mathbf{X}_x^* which minimizes the function M_x for a given dose-toxicity scenario x .

First, we can split the problem according to the therapeutic strategy s . For a given therapeutic strategy s , we denote by $S_s \subset \mathbb{R}^3$ the set of the authorized values, and we will look for:

$$\mathbf{X}_{s,x}^* = \underset{\mathbf{X} \in S_s}{\operatorname{argmin}} M_x(\mathbf{X}) \quad (\text{E.2})$$

$\mathbf{X}_{s,x}^*$ corresponds to the optimal parameter vector of strategy s , for the dose-toxicity scenario x . Previously, we stated that the constant strategy would be a special case of the decreasing or periodic ones if there were no constraints on the parameters. Thus, we impose to have $S_i \cap S_j = \emptyset$ for $1 \leq i < j \leq 3$. We consider S_s as a discrete set of values rather than a continuous subspace of \mathbb{R}^3 , so the optimization problem to solve will consist of looking at all potential values $\mathbf{X} \in S_s$ to find the one which minimizes $M_x(\mathbf{X})$. We refer to this procedure as "grid search". It would have been feasible to solve the optimization problem numerically even if S_s were continuous. Considering S_s as a discrete set of values will be convenient for looking for a trade-off strategy, as explained in the next section. Moreover, the parameter vector \mathbf{X} characterizes a therapeutic strategy; it would not be relevant to explore all potential values since the clinician could not apply exactly the dose strategy as we would recommend it. For example, it would not make sense for the clinician to have a precision of the day for the period L , or a precision of $1\mu\text{g}/\text{week}$ (i.e., $1/180$ for the normalized dose) for the choice of the dose \bar{d} . We consider the following potential values:

$$\bar{d} = (d_{inf}, d_{inf} + 0.02, d_{inf} + 0.04, \dots, 0.98, 1.0)$$

$$\mathbf{L} = (30, 60, \dots, 24 \times 30)$$

$$\boldsymbol{\lambda} = (0.05, 0.1, \dots, 0.90, 0.95)$$

To note that, already with the choice of \bar{d} , we consider more values than can be used in clinical routine. d_{inf} is specific to the patient, that is, a function of $\boldsymbol{\theta}$, as presented in § C.2.4. The period L is in days. We have $S_3 = \bar{d} \times \mathbf{L} \times \boldsymbol{\lambda}$, $S_2 = \bar{d} \times \mathbf{L} \times \{1.0\}$, and $S_1 = \bar{d} \times \{+\infty\} \times \{1.0\}$. For this latter, the singlet $\{+\infty\}$ could be replaced by any high enough value, for example, a time length that exceeds the maximal possible lifetime.

Once we have found the parameter vector which minimizes M_x for a given dose-toxicity scenario x and a given therapeutics strategy s , we can also look for the best strategy over the three considered. That is, we look for the following:

$$\mathbf{X}_x^* = \underset{\mathbf{X} \in S}{\operatorname{argmin}} M_x(\mathbf{X}) \quad (\text{E.3})$$

where $S = S_1 \cup S_2 \cup S_3$. The corresponding optimal strategy index s is such that $\mathbf{X}_x^* \in S_s$.

In this section, we only wanted to minimize the toxicity-related amount of IFN α given the dose-toxicity scenario x . When this latter is not known in advance, we have to search for a trade-off strategy.

E.2 Trade-off strategy

The best strategy for one dose-toxicity scenario is unlikely to be the best one for another scenario. Without prior knowledge of how IFN α toxicity increases with the dose, we must look for a compromise. Values of the toxicity-related amount of IFN α , as given by (E.1), are not comparable for two different toxicity relations. We consider S the set of all possible values for \mathbf{X} , over the three therapeutic strategies we consider. We denote by $|S| \in \mathbb{N}$ its cardinal, that is, the number of all potential values. Each value $\mathbf{X}_i \in S$, for $1 \leq i \leq |S|$, can be sorted according to the value $M_x(\mathbf{X}_i)$, for a given $1 \leq x \leq 4$. We note $n_x(\mathbf{X})$ the rank, with $n_x = 1$ being the best rank, and $q_x(\mathbf{X}) = n_x(\mathbf{X})/|S| \in [0, 1]$ (also expressed in % when presenting the results). We note $q(\mathbf{X}) = \max_{1 \leq x \leq 4} (q_x(\mathbf{X}))$. Then, finding the trade-off strategy is equivalent to finding:

$$\mathbf{X}^* = \operatorname{argmin}_{\mathbf{X} \in S} q(\mathbf{X}) \quad (\text{E.4})$$

E.3 Note about the decreasing strategy

The decreasing strategy is formulated in the article by:

$$d(t) = \begin{cases} \bar{d} & \text{for } t \in [T, T + L], \\ \lambda d(t - L) & \text{for } t \in [T + L, \tau] \end{cases} \quad (\text{E.5})$$

This is a recursive formulation. An other way to express it without self-referring is as follows:

$$d(t) = \bar{d} \cdot \lambda^{n(t)} \text{ for } t \in [T, \tau]$$

with $n(t) = \lfloor \frac{t-T}{L} \rfloor$ (where $\lfloor x \rfloor$ gives the integer part of x).

E.4 Note about the composite dose-toxicity relation

The composite dose-toxicity is defined in the article as followed:

$$z(d) = \frac{5\sqrt{5}}{9} \sqrt{d - 0.1} \mathbb{1}_{d \in [0.1, 0.55]} + \frac{5}{4\sqrt{5}} \frac{1}{\sqrt{1-d}} \mathbb{1}_{d \in [0.55, 1]} \quad (\text{E.6})$$

Here, we will detail the construction of this relation, and particularly the choice of the numerical values involved in that relation.

The composite relation is a more complex relation than the linear, convex, or concave ones, with:

- first, a low threshold d_{low} under which there is no toxicity: $z(d) = 0$ for $d < d_{low}$,
- second, a steep increase leading to a plateau like level, when a second threshold d_{up} is reached. We choose to model this increase by a relation close to our concave relation, that is, $z(d) \propto \sqrt{d - d_{low}}$ for $d \in [d_{low}, d_{up}[$ (ensuring the continuity of the function for $d = d_{low}$ with $z(d_{low}) = 0$),
- and finally, an increase of the toxicity towards an infinite value at the maximal considered dose $d_{max} = 1$. Several functions could have been chosen to model this behavior. In addition to reaching an infinite value for $d = d_{max}$, we also want the normalization $\int_0^1 z(d) dd = 1$. Such normalization requires to be able to integrate a function which diverges at a finite value. For that reason, a usual choice of function would be to consider that $z(d) \propto \frac{1}{\sqrt{d_{max}-d}}$ for $d \in [d_{up}, d_{max}]$.

Then, the choice for the composite relation would be:

$$z(d) = \begin{cases} 0 & \text{if } d \in [0, d_{low}[\\ \zeta_1 \sqrt{\frac{d-d_{low}}{d_{up}-d_{low}}} & \text{if } d \in [d_{low}, d_{up}[\\ \zeta_2 \sqrt{\frac{d_{max}-d_{up}}{d_{max}-d}} & \text{if } d \in [d_{up}, d_{max}] \end{cases}$$

with ζ_1 , ζ_2 , d_{low} and d_{up} to be determined ($d_{max} = 1$). Previous equation already accounts for two conditions : 1) continuity at $d = d_{low}$, and 2) divergence at $d = d_{max}$.

In addition to both conditions, we want: 3) continuity at $d = d_{up}$, that is:

$$\zeta_1 = \zeta_2 := \zeta \quad (\text{E.7})$$

4) continuity of the derivative at $d = d_{up}$, which leads to:

$$\frac{1}{2(d_{up} - d_{low})} = \frac{1}{2(1 - d_{up})}$$

that is:

$$d_{up} = 0.5(1 + d_{low}) \quad (\text{E.8})$$

and finally, 5) normalization:

$$\int_0^1 z(d) dd = 1 \quad (\text{E.9})$$

We have:

$$\begin{aligned} \int_0^1 z(d) dd &= \int_0^{d_{low}} z(d) dd + \int_{d_{low}}^{d_{up}} z(d) dd + \int_{d_{up}}^1 z(d) dd \\ &= 0 + \frac{\zeta}{\sqrt{d_{up} - d_{low}}} \int_0^{d_{up} - d_{low}} \sqrt{d} dd + \zeta \sqrt{1 - d_{up}} \int_0^{1 - d_{up}} \frac{1}{\sqrt{d}} dd \\ &= \frac{\zeta}{\sqrt{d_{up} - d_{low}}} \frac{2}{3} (d_{up} - d_{low})^{3/2} + 2\zeta(1 - d_{up}) \\ &= \zeta \left(\frac{2}{3} (d_{up} - d_{low}) + 2(1 - d_{up}) \right) \end{aligned}$$

The previous conditions are not sufficient to allow determining a value of each our parameters; there is still one degree of liberty left. We choose to set $d_{low} = 0.1$ (corresponding to 18 $\mu\text{g}/\text{week}$).

Then, from eq. (E.8), we get $d_{up} = 0.55$ (that is, $d_{up} - d_{low} = 0.45 = 1 - d_{up}$).

Finally, $\int_0^1 z(d) dd = 0.45\zeta \left(\frac{2}{3} + 2 \right) = \frac{6}{5}\zeta$ so that $\zeta = \frac{5}{6}$ according to condition (E.9).

Using these numerical values, the composite dose-toxicity can be written as:

$$z(d) = \begin{cases} 0 & \text{if } d \in [0, 0.1[\\ \frac{5}{6} \sqrt{\frac{d-0.1}{9/20}} & \text{if } d \in [0.1, 0.55[\\ \frac{5}{6} \sqrt{\frac{9/20}{1-d}} & \text{if } d \in [0.55, 1] \end{cases}$$

or equivalently, $z(d)$ can be given by eq. (E.6).

F Estimating the best timing for measuring the clonal architecture

Our recommendations must be based on accurate estimations of the individual model parameters for validity. The clinician might play an active role not only in the prescription of the IFN α dose and the monitoring of the hematological response but also in the choice of the dates at which the clonal architecture of progenitor cells is measured. Indeed, by default, we chose to consider that the second measurement (the first one being at the start of the therapy) of the CF among immature cells is obtained about 300 days after the start of the therapy. This choice is consistent with another study exploring the effect of combining IFN α with ruxolitinib on MPN patients, where the clonal architecture could only be measured twice, and the second time-point was chosen about 300 days after the start of the therapy [10]. In this part, we want to explore the relevance of this choice.

To this end, we will conduct a study based on a synthetic dataset derived from the three patients #18, 12, and 32, for whom we estimated the model parameters.

For a given patient i , we estimated their parameter vector $\theta^{(i)}$ as presented in § C.2.2, here with the assimilation time $T = 600$ days. We assume these values to be the true ones. The parameter vector entirely describes the patient's response to the treatment for any potential dose $d(t) \in [0, 1]$ over time t . We simulate, from the model and the value for $\theta^{(i)}$, the response to the treatment for $t \in [0, 600]$ and the dose $d(t)$ received by the considered patient. We get the simulated dynamics of the VAF $y_{\theta^{(i)}}(t)$ and the heterozygous $z_{\theta^{(i)},het}(t)$ and homozygous $z_{\theta^{(i)},hom}(t)$ CF over time t . We call the latter values the true ones (or the theoretical ones). We will try to retrieve these values using our inference method. In particular, we will explore the influence of the choice of the time point T_{obs} of the second observed CF. For that purpose, we simulate synthetic datasets (Fig. F.1-i).

For the considered patient i , the (clinical) VAF was measured at different time points: $\left\{ t_k^{(i)} \right\}_{k \in \mathcal{I}_{T=600}^{(i)}}$.

We consider then the simulated values of the VAF $y_{\theta^{(i)}}(t_k^{(i)})$ for $k \in \mathcal{I}_{T=600}^{(i)}$. From these latter, we generate noisy VAF measurements $\hat{y}_k^{(i)}$ using eq. (A.1) (pseudo-observations). Then, we build the following synthetic datasets, according to the choice of the time $T_{obs} \in \{1, \dots, 600\}$, which corresponds to the time when the second (pseudo-)observation of the clonal architecture would be made (Fig. F.1-ii):

$$\mathcal{D}_{T_{obs}}^{(i)} = \left\{ t_k^{(i)}, \hat{y}_k^{(i)} \right\}_{k \in \mathcal{I}_{T=600}^{(i)}} \cup \left\{ z_{\theta^{(i)},het}(0), z_{\theta^{(i)},hom}(0) \right\} \cup \left\{ z_{\theta^{(i)},het}(T_{obs}), z_{\theta^{(i)},hom}(T_{obs}) \right\}$$

We conduct 600 parameter estimation procedures (Fig. F.1-iii), one for each $\mathcal{D}_{T_{obs}}^{(i)}$, $1 \leq T_{obs} \leq 600$, to estimate the parameter vector $\theta_{T_{obs}}^{(i)}$ for which the posterior density function is maximal. For each $T_{obs} \in \{1, \dots, 600\}$, we estimate the parameter $\hat{\theta}_{T_{obs}}^{(i)}$ which maximizes the posterior:

$$p[\theta_{T_{obs}}^{(i)} | \mathcal{D}_{T_{obs}}^{(i)}] \propto p[\mathcal{D}_{T_{obs}}^{(i)} | \theta_{T_{obs}}^{(i)}] p[\theta_{T_{obs}}^{(i)}] \quad (\text{F.1})$$

with $p[\theta_{T_{obs}}^{(i)}]$ the prior expressed in § C.2.1.

Since the optimization problem has to be repeated 600 times for each patient, we solve numerically with an efficient algorithm, namely the CMA-ES (Covariance Matrix Adaptation - Evolution Strategy) algorithm [4, 5]. $\hat{\theta}_{T_{obs}}^{(i)}$ is estimated numerically using the CMA-ES algorithm. From our model and the value of $\hat{\theta}_{T_{obs}}^{(i)}$, we get the inferred dynamics of the heterozygous $z_{\hat{\theta}_{T_{obs}}^{(i)},het}(t)$ and homozygous $z_{\hat{\theta}_{T_{obs}}^{(i)},hom}(t)$ CF, that we can confront to the theoretical ones, that is, to the target ones obtained by simulation (Fig. F.1-iv). For that purpose, we define the following prediction error:

$$\begin{aligned} \varepsilon(T_{obs}) &= \int_0^{T=600} |z_{\hat{\theta}_{T_{obs}}^{(i)},het}(t) - z_{\theta^{(i)},het}(t)| + |z_{\hat{\theta}_{T_{obs}}^{(i)},hom}(t) - z_{\theta^{(i)},hom}(t)| dt \\ &\approx \sum_{k=1}^{600} |z_{\hat{\theta}_{T_{obs}}^{(i)},het}(k) - z_{\theta^{(i)},het}(k)| + |z_{\hat{\theta}_{T_{obs}}^{(i)},hom}(k) - z_{\theta^{(i)},hom}(k)| \end{aligned} \quad (\text{F.2})$$

Then, we search for the (pseudo-)observation time T_{obs}^* which minimizes the previous error:

$$T_{obs}^* = \underset{1 \leq T_{obs} \leq 600}{\operatorname{argmin}} \varepsilon(T_{obs}) \quad (\text{F.3})$$

From this study, we could then discuss whether there would be some choices that are better than others (and consistent among the three patients we consider), to deduce the best timing for the second measurement of clonal architecture.




Figure F.1: Illustration of the experimental design method. We consider that we know the true parameter vector for a given patient, that is, the parameter vector θ . First, we simulate data (i), consisting of VAF measurements, as well as daily observations of the clonal architecture (pseudo-observations). Note that we simulate as many VAF pseudo-observations as we actually have for the patient and at the same observation times. For clarity, we do not represent all the data points for the CF. We consider different potential pseudo-datasets for the parameter estimation procedure; these pseudo-datasets are built with all VAF pseudo-observations and only two pseudo-observations of the CF (ii). The pseudo-datasets differ according to the choice of the second pseudo-observation time T_{obs} of the CF, here t_1 vs t_2 . For each pseudo-dataset, we infer the model parameters, and, thus, the cell dynamics over IFN α (iii). Then, we compare the error between the inferred CF dynamics and the true one (iv).

G Detailed results

Our study is organized into three parts, and so are our results. In § G.1, we present the results of the prediction, in § G.2 how the therapy could be optimized, and in § G.3 how to choose the best timing for measuring the clonal architecture. We present the results separately for patients #12, 18, and 32.

G.1 Prediction

G.1.1 Patient #12

We exclude patient #12 from the cohort of the 19 $JAK2^{V617F}$ patients. For the 18 remaining patients, we run a hierarchical Bayesian estimation, as presented in § C.1. For patient #12, we will estimate their parameters, as presented in § C.2, for different assimilation times T . According to the values of T , a different number of VAF measurements are used for the estimation:

- 5 VAF measurements for $T = 300$ days,
- 7 VAF measurements for $T = 600$ days,
- and 10 VAF measurements for $T = 1,000$ days.

In all cases, two measurements of the CF among progenitors are considered, the first at $t = 0$ and the second at $t = 287$ days.

The posterior distributions of the parameters are presented in Fig. G.2, and the results of the predictions are presented in Fig. G.1. The MSE are presented in Tab. G.1. Note that, for $t \geq 1,000$ days, we do not have measurements of the VAF anymore, so we can not compute the prediction error MSE_{pred}^{VAF} .

We get a good agreement between the observed and predicted VAF values already for $T = 300$ days, but we underestimate the values of the homozygous CF. The credibility intervals are large for $T = 300$ days and decrease for $T = 600$ and $T = 1,000$ days.

When looking at the posterior distributions, we can see the impact of the prior knowledge learned from the cohort: the posterior distributions - in particular those of parameters related to heterozygous cells (γ_{het} , $\delta_{0,h}$, ρ_{het}) - are close to the prior distributions, especially for $T = 300$. For η_{het} and η_{hom} , the prior is chosen uniform and, thus, does not come from the cohort. For $\delta_{0,h}$ and γ_{hom} , their posteriors for $T = 300$ differ from the prior, and we can see how increasing T updates the posteriors. Adding more VAF measurements also seems to impact the estimation of ρ_{hom} ; the more data, the less variance we have for the posterior.

However, most of the information comes from the observations before $T = 300$ (and from the prior): the posterior we obtain with $T = 600$ or 1,000 does not differ much from the one we get with $T = 300$. The reason is that, when increasing T , we only add VAF measurements but no additional observations of the CF among progenitors. In general, the posterior distributions of patient #12 are close to the prior, suggesting that their response to the therapy does not deviate too large from those obtained for the 18

	$T = 300$			$T = 600$			$T = 1000$		
	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$
VAF	0.0031	0.0011	0.0021	0.0016	0.0015	0.0015	0.0011	NA	0.001
het CF	0.0087	0.0061	0.0067	0.0108	0.0080	0.0086	0.0111	0.0083	0.0089
hom CF	0.0257	0.0377	0.0353	0.0282	0.0417	0.0390	0.0298	0.0442	0.0414
All	0.0375	0.0449	0.0440	0.0406	0.0512	0.0491	0.0420	NA	0.0514

Table G.1: Results for patient # 12. Values of the MSE computed for different assimilation times T (300, 600, or 1000 days), different types of measurements (VAF among mature cells, heterozygous or homozygous CF among progenitors, or all together), and different observations: either from the control dataset ($\mathcal{D}_c^{(i)}$), the one used for the model calibration ($\mathcal{D}_T^{(i)}$), or the whole dataset ($\mathcal{D}^{(i)}$). For $t \geq 1,000$ days, we do not have measurements of the VAF, so we can not compute the prediction error and displays NA (Not applicable) instead.

remaining patients in the cohort. However, we can also see that there is not a good agreement between the observed and inferred homozygous CF at $t = 287$, suggesting that the dynamics of the VAF would agree with what could be expected for the other patients but not the dynamic of the CF. More broadly, when looking at the residuals, we observe systematic deviations for the heterozygous (overestimation) and homozygous (underestimation) values (most of them being not used for estimating the model parameters). This is in line with the previous comment suggesting that the observed dynamics among progenitor cells for patient # 12 might differ from what would be expected *a priori*. For more reliable predictions concerning the heterozygous and homozygous CF dynamics, more than two measurements of the heterozygous and homozygous CF should be considered when fitting the model.

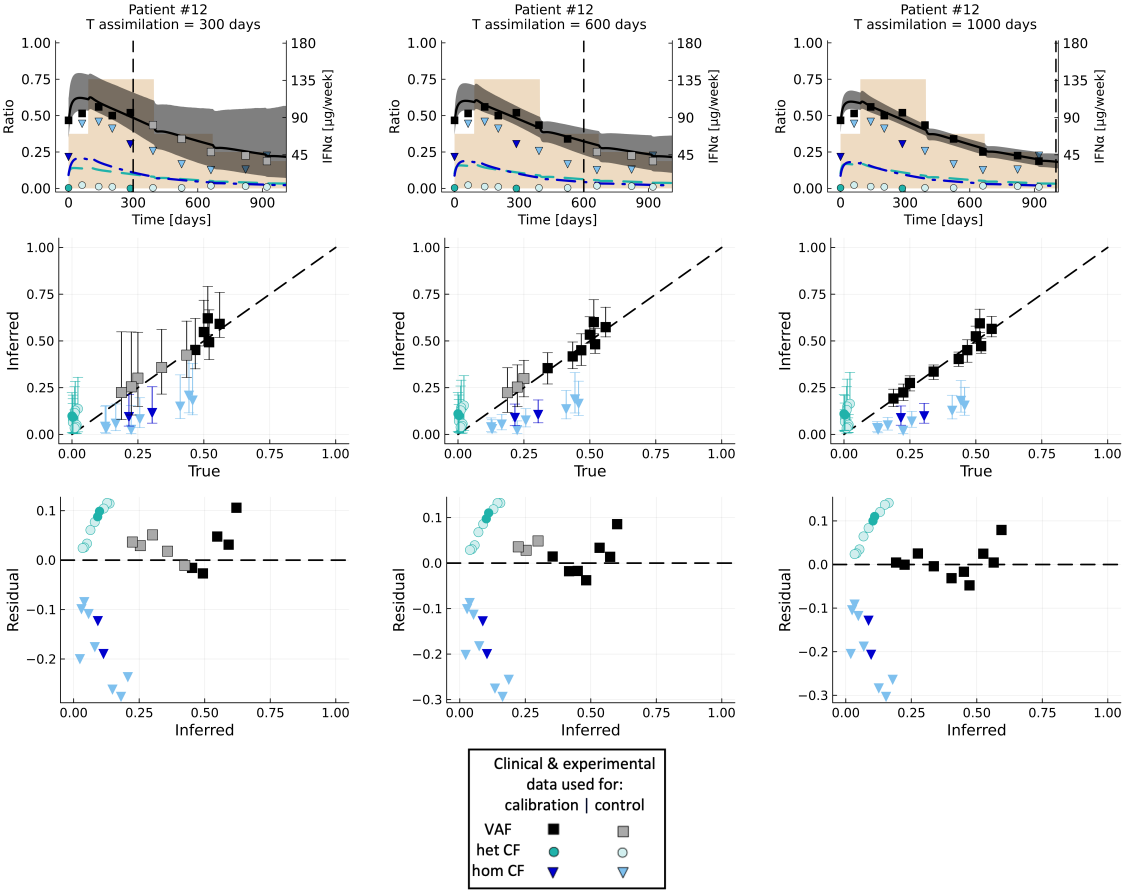


Figure G.1: Data assimilation results for patient #12. From left to right, the assimilation times are $T = 300$, 600, and 1,000 days (vertical dash line). At the top, we present the predicted dynamics. In black, the VAF (median value and 95% credibility interval), and respectively in green and blue, the CF (medians) of heterozygous and homozygous mutated progenitors. The squares, circles, and triangles correspond to the clinical/experimental data, either used to calibrate the model (dark colors), or to control the quality of the predictions (light colors - control dataset). In brown, we represent the IFN α dose variations over treatment (in $\mu\text{g}/\text{week}$). Second line, we compare the inferred values (median values, in the y-axis) with those observed (in the x-axis) for both the control data set and the one used for the model calibration. The error bars correspond to the 95% credibility intervals. At the bottom, we display the residuals (inferred values minus the observed ones) (y-axis) according to the inferred values (x-axis).

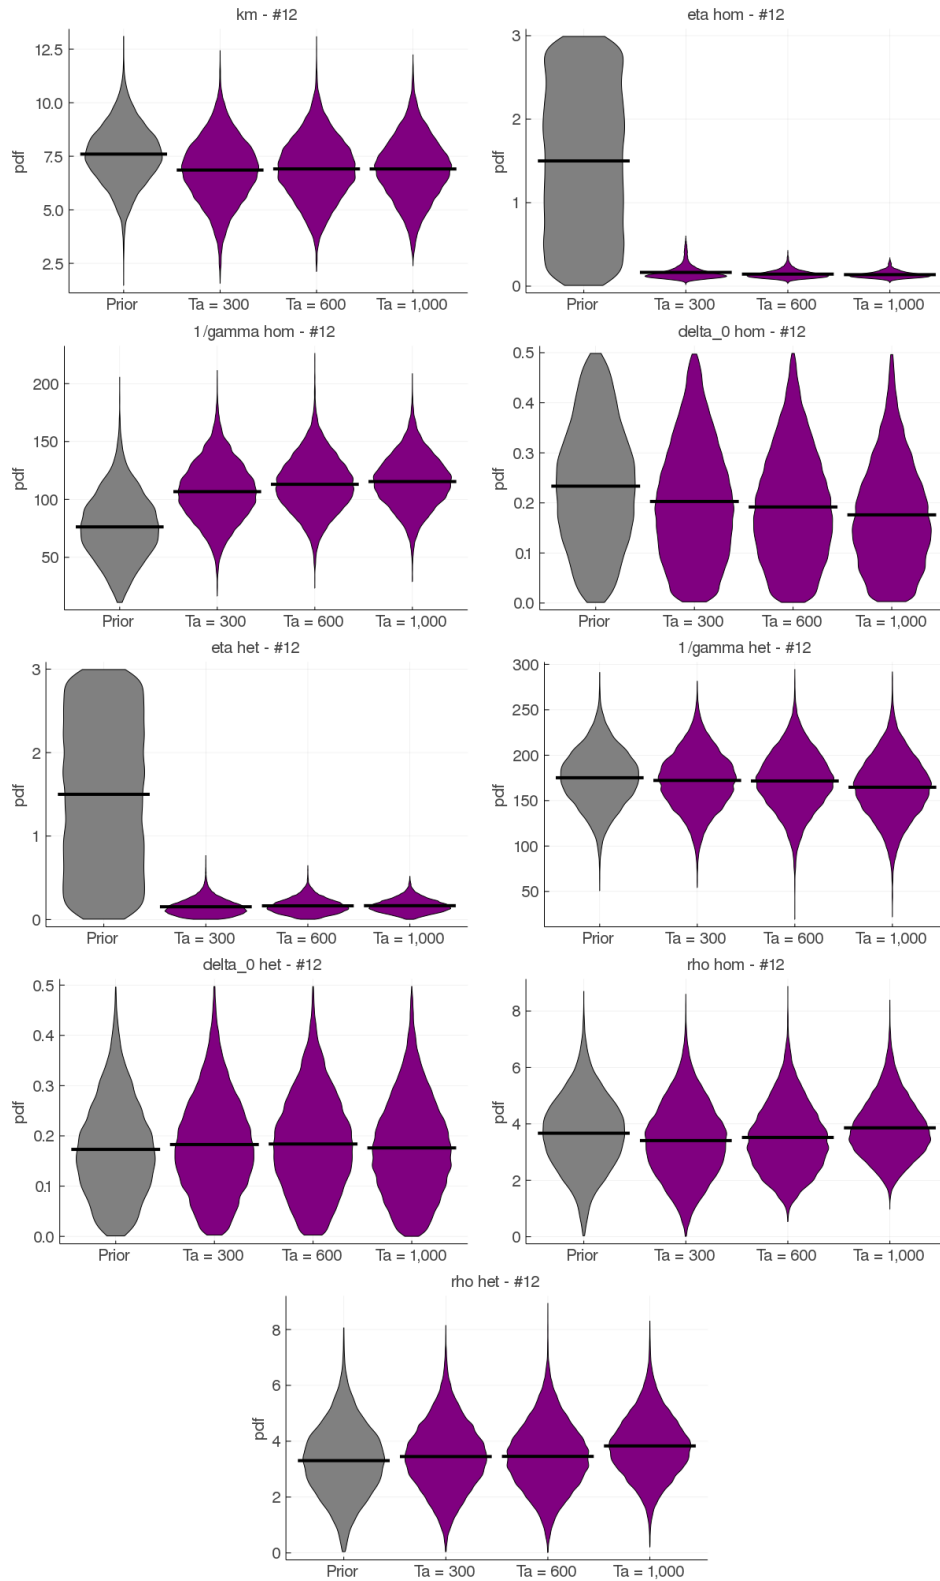


Figure G.2: Posterior distributions (pdf - in purple) of all parameters, for patient #12, according to the assimilation time T . In grey on the left, we display the prior distribution (that is, the one obtained from the hierarchical inference based on the 18 remaining patients). For parameters η_{het} and η_{hom} , which refer to the initial conditions, we choose a non-informative uniform prior. Otherwise, the prior comes from the hierarchical Bayesian estimation made with the data of the 18 remaining patients. The horizontal dark line represents the mean value.

G.1.2 Patient #18

We exclude patient #18 from the cohort of the 19 $JAK2^{V617F}$ patients. For the 18 remaining patients, we run a hierarchical Bayesian estimation, as presented in § C.1. For patient #18, we will estimate their parameters, as presented in § C.2, for different assimilation times T . According to the values of T , a different number of VAF measurements are used for the estimation:

- 4 VAF measurements for $T = 300$ days,
- 8 VAF measurements for $T = 600$ days,
- and 11 VAF measurements for $T = 1,000$ days.

In all cases, two measurements of the CF among progenitors are considered, the first at $t = 0$ and the second at $t = 248$ days.

The parameter posterior distributions are presented in Fig. G.4 and the results of the predictions are presented in Fig. G.3. The MSE are presented in Tab. G.2.

The results for patient #18 were detailed in §3.1 in the main text. In general, we can see a good agreement between the inferred and observed values, both for the VAF and the heterozygous and homozygous CF. We did not observe systematic deviations concerning the observations used for the model calibration. However, we tend to underestimate (for $T = 600$ or $T = 1000$) the predicted values of the homozygous CF when overestimating those of the heterozygous CF, such that there is a compensation leading to a good estimation of the VAF.

For this patient, it is interesting to see how the posterior distribution of some parameters ($\delta_{0,hom}$, $\delta_{0,het}$, ρ_{hom} , γ_{hom}) is updated when adding more and more data (i.e., when increasing the assimilation time T) when for other parameters, the posterior stays close from the prior (k_m , γ_{het}). In particular, the posterior of $\delta_{0,hom}$ largely deviates from the prior, already when $T = 300$, and even more for higher T . It is also interesting to observe that the posterior distributions of η_{het} and η_{hom} - parameters associated with the initial quantities of heterozygous and homozygous mutated cells - remain unchanged when increasing the assimilation time.

	$T = 300$			$T = 600$			$T = 1000$		
	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$
VAF	0.0068	0.0066	0.0067	0.0040	0.0006	0.0026	0.0026	0.0005	0.0021
het CF	0.0002	0.0047	0.0041	0.0016	0.0085	0.0076	0.0027	0.0163	0.0145
hom CF	0.0086	0.0065	0.0068	0.0046	0.0045	0.0045	0.0034	0.0062	0.0058
All	0.0155	0.0177	0.0175	0.0102	0.0136	0.0147	0.0087	0.0231	0.0225

Table G.2: Results for patient # 18. Values of the MSE computed for different assimilation times T (300, 600, or 1000 days), different types of measurements (VAF among mature cells, heterozygous or homozygous CF among progenitors, or all together), and different observations: either from the control dataset ($\mathcal{D}_c^{(i)}$), the one used for the model calibration ($\mathcal{D}_T^{(i)}$), or the whole dataset ($\mathcal{D}^{(i)}$).

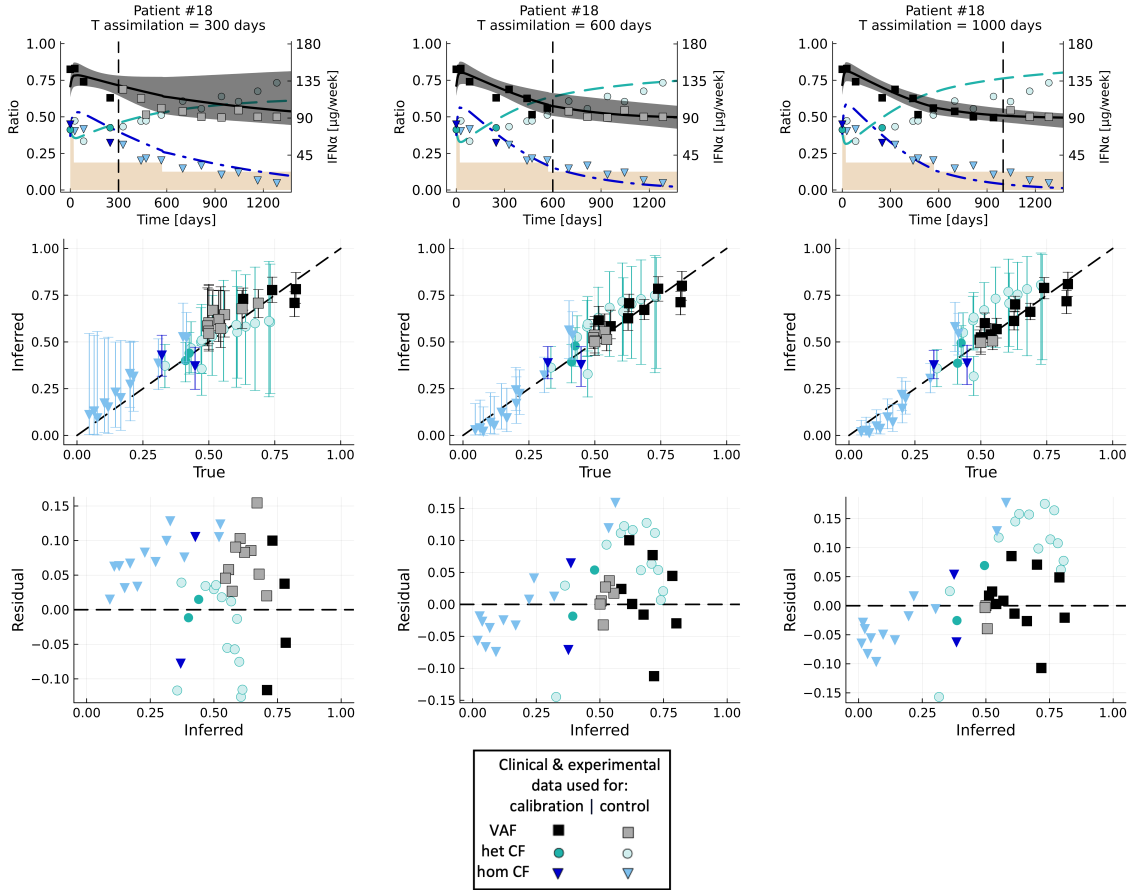


Figure G.3: Data assimilation results for patient #18. From left to right, the assimilation times are $T = 300$, 600, and 1,000 days (vertical dash line). At the top, we present the predicted dynamics. In black, the VAF (median value and 95% credibility interval), and respectively in green and blue, the CF (medians) of heterozygous and homozygous mutated progenitors. The squares, circles, and triangles correspond to the experimental/clinical data, either used to calibrate the model (dark colors), or to control the quality of the predictions (light colors - control dataset). In brown, we represent the IFN α dose variations over treatment (in $\mu\text{g}/\text{week}$).

Second line, we compare the inferred values (median values, in the y-axis) with those observed (in the x-axis) for both the control data set and the one used for the model calibration. The error bars correspond to the 95% credibility intervals. At the bottom, we display the residuals (inferred values minus the observed ones) (y-axis) according to the inferred values (x-axis).

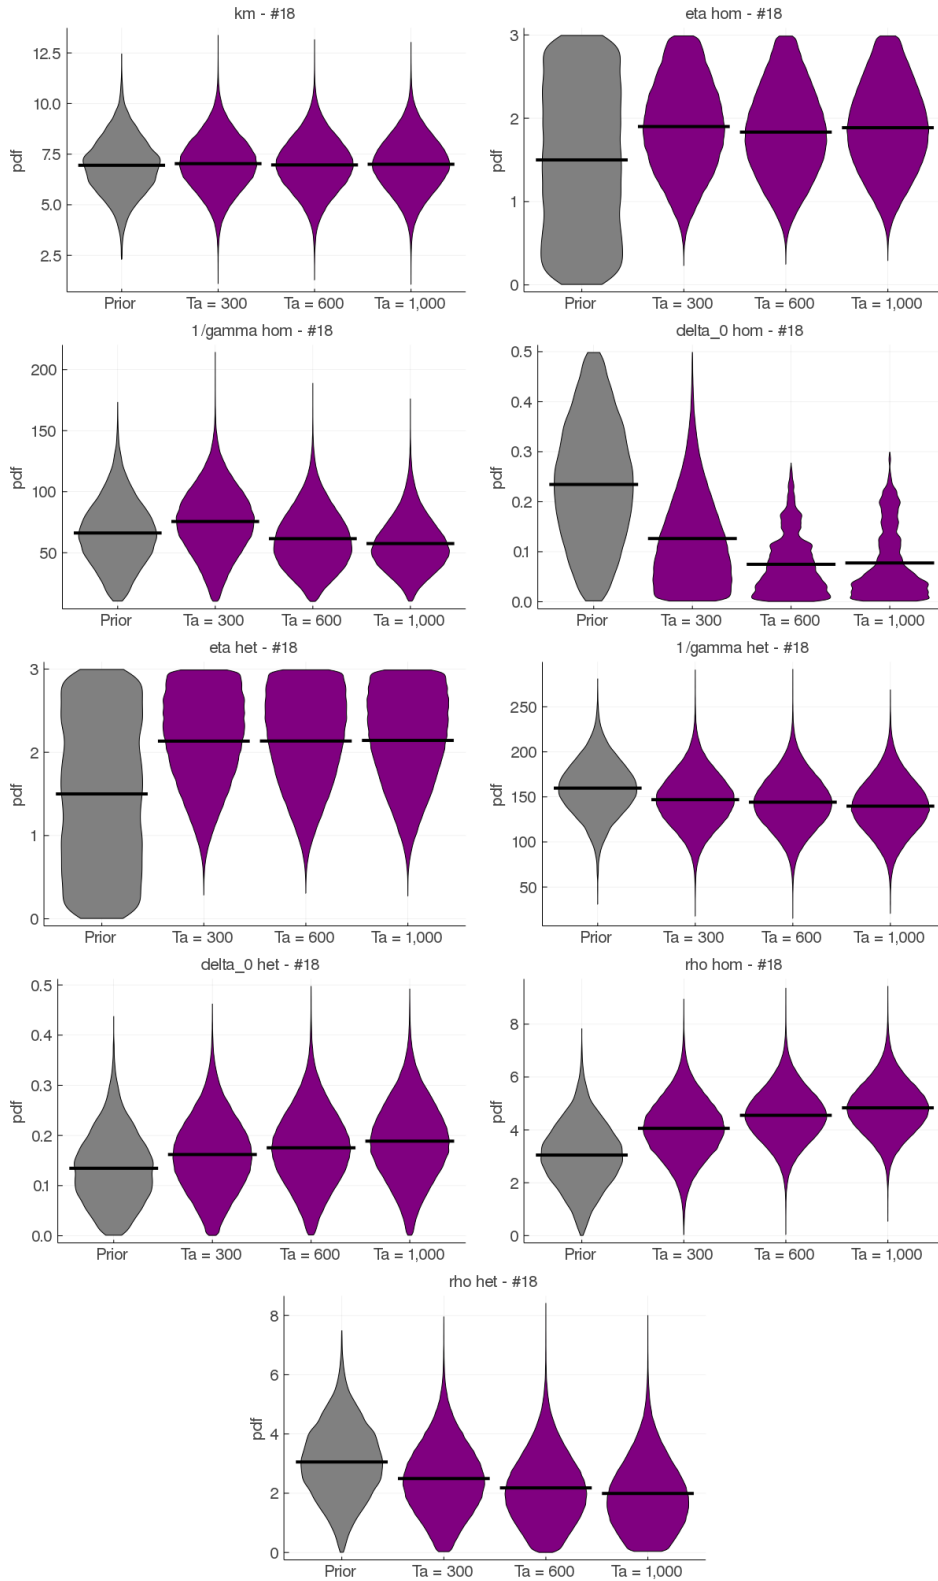


Figure G.4: Posterior distributions (pdf - in purple) of all parameters, for patient #18, according to the assimilation time T . In grey on the left, we display the prior distribution. For parameters η_{het} and η_{hom} , which refer to the initial conditions, we choose a non-informative uniform prior. Otherwise, the prior comes from the hierarchical Bayesian estimation made with the data of the 18 remaining patients. The horizontal dark line represents the mean value.

G.1.3 Patient #32

We exclude patient #32 from the cohort of the 19 $JAK2^{V617F}$ patients. For the 18 remaining patients, we run a hierarchical Bayesian estimation, as presented in § C.1. For patient #32, we will estimate their parameters, as presented in § C.2, for different assimilation times T . According to the values of T , a different number of VAF measurements are used for the estimation:

- 3 VAF measurements for $T = 300$ days,
- 5 VAF measurements for $T = 600$ days,
- and 8 VAF measurements for $T = 1,000$ days.

In all cases, two measurements of the CF among progenitors are considered, the first at $t = 0$ and the second at $t = 235$ days.

The results of the predictions are presented in Fig. G.5. The MSE are presented in Tab. G.3. For $T = 600$, we get poor predictions. Therefore, we study the impact of the choice of another time point for measuring the clonal architecture among progenitors. The results, presented in the next paragraph, show that the poor predictions we get for $T = 600$ can be explained by the choice of the CF used for inferring the dynamics. In other words, the time point for measuring the clonal architecture matters.

	$T = 300$			$T = 600$			$T = 1000$		
	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$
VAF	0.0164	0.0217	0.0205	0.0049	0.2500	0.1557	0.0065	0.0235	0.0130
het CF	0.0022	0.0010	0.0012	0.0033	0.0007	0.0011	0.0012	0.0010	0.0010
hom CF	0.0073	0.0060	0.0062	0.0006	0.0862	0.0730	0.0015	0.0073	0.0064
All	0.0164	0.0287	0.0279	0.0088	0.3370	0.2299	0.0065	0.0318	0.0205

Table G.3: Results for patient # 32. Values of the MSE computed for different assimilation times T (300, 600, or 1000 days), different types of measurements (VAF among mature cells, heterozygous or homozygous CF among progenitors, or all together), and different observations: either from the control dataset ($\mathcal{D}_c^{(i)}$), the one used for the model calibration ($\mathcal{D}_T^{(i)}$), or the whole dataset ($\mathcal{D}^{(i)}$).

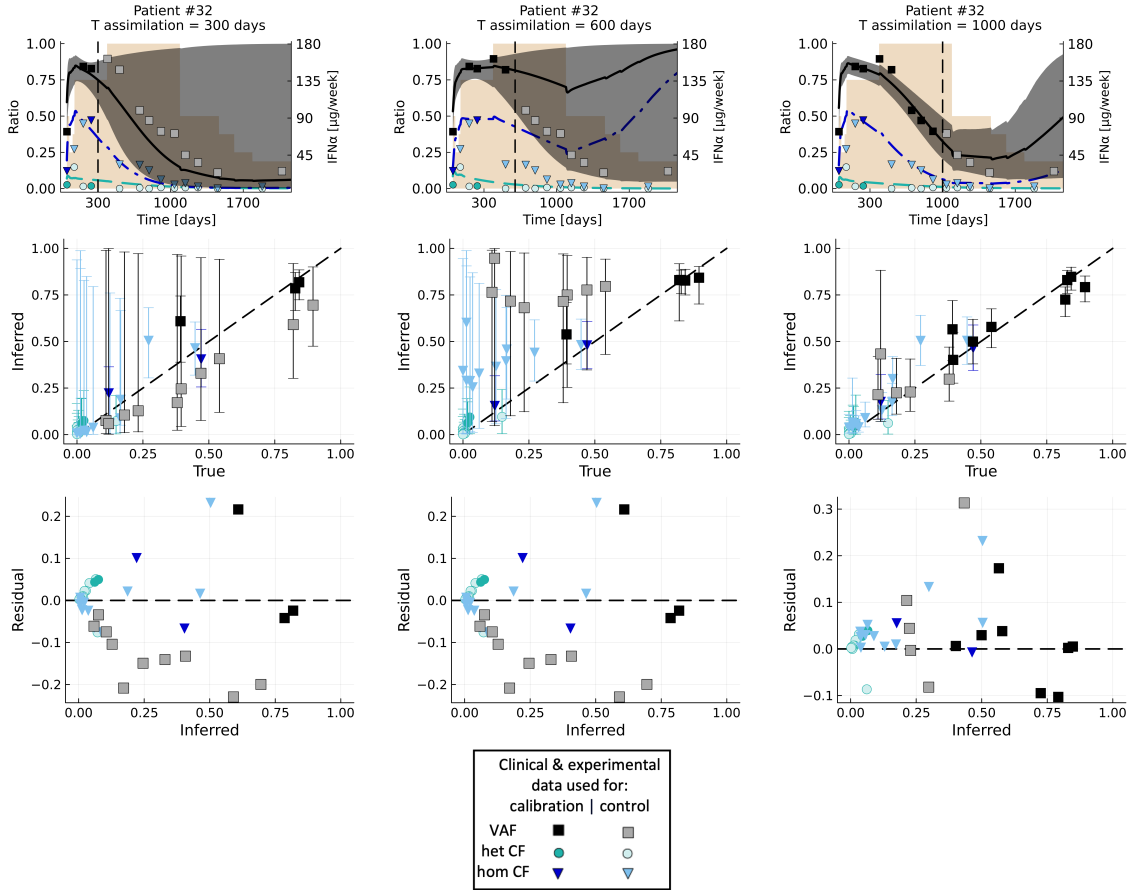


Figure G.5: Data assimilation results for patient #32. From left to right, the assimilation times are $T = 300$, 600, and 1,000 days (vertical dash line). At the top, we present the predicted dynamics. In black, the VAF (median value and 95% credibility interval), and respectively in green and blue, the CF (medians) of heterozygous and homozygous mutated progenitors. The squares, circles, and triangles correspond to the clinical/experimental data, either used to calibrate the model (dark colors), or to control the quality of the predictions (light colors - control dataset). In brown, we represent the IFN α dose variations over treatment (in μ g/week).

Second line, we compare the inferred values (median values, in the y-axis) with those observed (in the x-axis) for both the control data set and the one used for the model calibration. The error bars correspond to the 95% credibility intervals. At the bottom, we display the residuals (inferred values minus the observed ones) (y-axis) according to the inferred values (x-axis).

G.1.4 Patient #32 when choosing another time point for measuring the clonal architecture

As we saw in the previous paragraph, the predictions were not good for patient #32 for $T = 600$ and a CF measured at $t = 235$ days. If we run the same computation, but instead of choosing the second CF observation at $t = 235$ days, we consider the one at $t = 508$ days, we obtain the results presented in Fig. G.6. The predictions for $T = 600$ are much better (Tab. G.4), illustrating the importance of the choice of the time point for measuring the clonal architecture. No systematic deviations are observed. When looking at the posterior distributions (Fig. G.7), we can observe that the posterior distributions of the parameters related to heterozygous mutated cells do not deviate from the prior, which is explained by the fact that patient #32 has almost no heterozygous cells. There is especially one parameter - namely ρ_{hom} - whose posterior distribution is updated when increasing the number of VAF measurements.

	$T = 600$			$T = 1000$		
	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$	$\mathcal{D}_T^{(i)}$	$\mathcal{D}_c^{(i)}$	$\mathcal{D}^{(i)}$
VAF	0.0133	0.0075	0.0097	0.0071	0.0161	0.0106
het CF	0.0006	0.0010	0.0009	0.0004	0.0011	0.0010
hom CF	0.0063	0.0038	0.0042	0.0051	0.0038	0.0040
All	0.0201	0.0123	0.0148	0.0126	0.0210	0.0156

Table G.4: Results for patient # 32, when considering a measurement of the clonal architecture at $t = 0$ and $t = 508$ (Therefore, we do not present the results for $T = 300$). Values of the MSE computed for different assimilation times T (600 or 1000 days), different types of measurements (VAF among mature cells, heterozygous or homozygous CF among progenitors, or all together), and different observations: either from the control dataset ($\mathcal{D}_c^{(i)}$), the one used for the model calibration ($\mathcal{D}_T^{(i)}$), or the whole dataset ($\mathcal{D}^{(i)}$).




Figure G.6: Data assimilation results for patient #32, when considering a measurement of the clonal architecture at $t = 0$ and $t = 508$ (Therefore, we do not display the results for $T = 300$). On the left, we have $T = 600$, and on the right, $T = 1,000$. At the top, we present the predicted dynamics. In black, the VAF (median value and 95% credibility interval), and respectively in green and blue, the CF (medians) of heterozygous and homozygous mutated progenitors. The squares, circles, and triangles correspond to the clinical/experimental data, either used to calibrate the model (dark colors), or to control the quality of the predictions (light colors - control dataset). In brown, we represent the IFN α dose variations over treatment (in μ g/week). Second line, we compare the inferred values (median values, in the y-axis) with those observed (in the x-axis) for both the control data set and the one used for the model calibration. The error bars correspond to the 95% credibility intervals. At the bottom, we display the residuals (inferred values minus the observed ones) (y-axis) according to the inferred values (x-axis).




Figure G.7: Posterior distributions (pdf - in purple) of all parameters, for patient #32, according to the assimilation time T . Here, we consider that the CF is measured at $t = 0$ and $t = 508$ (instead of 235). In grey on the left, we display the prior distribution. For parameters η_{het} and η_{hom} , which refer to the initial conditions, we choose a non-informative uniform prior. Otherwise, the prior comes from the hierarchical Bayesian estimation made with the data of the 18 remaining patients. The horizontal dark line represents the mean value.

G.1.5 Synthesis

This section summarises the results we get for patients #12, 18, and 32. In Tab. G.5, we display the predicted error based on the VAF values (MSE_{pred}^{VAF}). In Tab. G.6, we present the estimated parameter vector (mean posterior value) for the assimilation time $T = 600$. These are the values used when studying the treatment optimization (section G.2) and the experimental design (section G.3). The parameter k_m is similar for each of our three patients. It might suggest that the proliferative advantage of the $JAK2^{V617F}$ mutated cells at the last stages of hematopoiesis is a property of the mutated cell and not patient-dependent. η_{het} and η_{hom} correspond to the estimated initial quantities of heterozygous and homozygous mutated HSCs compared to the quantities of WT HSCs. Patient #18 initially has far more mutated cells than the two others, which explains why patient #18 is diagnosed with a PMF, a more advanced disease than PV. The values of $1/\gamma_{het}$ do not much differ between the three patients, and actually, do not differ much from the prior (that is, from the population distribution), suggesting that the effect of $\text{IFN}\alpha$ on the quiescence exit of mutated HSCs does not differ much between patients. For homozygous mutated HSCs, on the contrary, the value of $1/\gamma_{hom}$ differs according to the patient, with a trend to have lower values (that is, higher rate of quiescence exit) when there are more mutated HSCs (but three patients is too few to have a statistical significance). Overall, patients # 12 and #18, who have few heterozygous mutated cells, have similar estimated parameter values when associated with heterozygous mutated cells. This is explained by the fact that, for them, the posterior distributions of the parameter associated with the heterozygosity will not differ much from the prior. Patient # 18 will have their heterozygous mutated HSCs slightly less targeted than the two other patients (according to the estimated value for ρ_{het}). However, on the other hand, their homozygous mutated cells will be targeted more efficiently (since patient # 18 has a lower value for $\delta_{0,hom}$ and a higher value for ρ_{hom} , compared to the two others). Here again, the potential difference between patient #18 *vs* patients #12 and 32 might lie in the difference at the beginning of the therapy, that is, a much higher mutated heterozygous and homozygous CF before the start of the treatment.

To sum up, differences between the estimated model parameters could be, of course, explained by an inter-individual heterogeneity, but also by the fact that the patients have different prognostic when they start the therapy, with patient # 18 having the most severe one (PMF) associated with a high proportion of both heterozygous and homozygous mutated cells.

T	#12	#18	#32	#32 (bis)
300	0.0011	0.0066	0.0217	NR
600	0.0015	0.0006	0.2500	0.0075
1,000	NR	0.0005	0.0235	0.0161

Table G.5: Results of the prediction error MSE_{pred}^{VAF} . NR means Not Relevant. #32 (bis) corresponds to the dataset of patient #32 when choosing the second observation for the CF at $t = 508 > 300$.

We also explored the relevance of our observation model - presented in eq. (A.1) - with a QQ-plot in Fig. G.8 or with the residual plots presented at the bottom of Fig. G.1, G.3, G.5, and G.6. However, these results suggest that the assumption of Gaussian distributed errors might not be totally relevant, and other choices could be more appropriate. It should also be recalled that the dynamic model used to describe the on-treatment dynamics of $JAK2^{V617F}$ patients is necessarily a simplification of reality. Consequently, the model could also introduce some biases, explaining why the residuals are not randomly distributed around zero in Fig. G.1, G.3, G.5, and G.6.

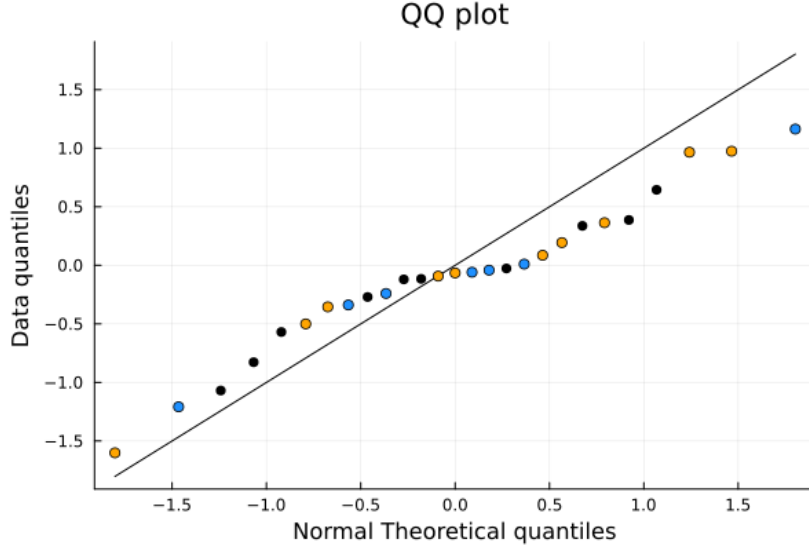


Figure G.8: QQ-plot showing to which extent the assumption of Gaussian distributed errors, as expressed by equation (A.1), holds for the VAF observations (when used to estimate the model parameters). According to equation (A.1), $\hat{y}_k^{(i)} | y_k^{(i)} \sim \mathcal{N}\left(y_k^{(i)}, y_k^{(i)}(1 - y_k^{(i)})\sigma_m^2\right)$, that is, $\frac{\hat{y}_k^{(i)} - y_k^{(i)}}{\sigma_m \sqrt{y_k^{(i)}(1 - y_k^{(i)})}}$ (y-axis) should be distributed

according to a standard Gaussian law (x-axis). Here, $\hat{y}_k^{(i)}$ corresponds to the VAF observed at time $t_k^{(i)}$ for individual i , when $y_k^{(i)}$ corresponds to the theoretical value (median values obtained after propagating the uncertainties from the posterior distribution of the parameters to the model output - computations made with the assimilation time $T = 1,000$ days). The values used for computing the QQ-plot are associated with the VAF measurements obtained before $T = 1,000$ days of treatment for patient #12 (orange circles), #18 (black), and #32 (blue).

The QQ-plot differs from what would be ideally expected if the errors were truly distributed according to eq. (A.1) and if the model was the real one. Our model is necessarily a simplification of reality; it is not excluded that it has some bias and might not be totally suitable to describe the clinical data. However, if not perfect, the considered model resulted from a model selection procedure performed in [9] and therefore turned out to be the most appropriate given a large set of (217) potential models. More likely, the assumption of Gaussian distributed errors, as we initially introduced it in [15], might not be verified, and it would be relevant to explore if some noise models were more suitable.

Parameter	#12	#18	#32 (bis)	Baseline (WT)
k_m	6.97	6.96	7.22	1.0
η_{hom}	0.14	1.83	0.36	-
$1/\gamma_{hom}$	113.0	61.58	80.4	300.0
$\delta_{0,hom}$	0.19	0.07	0.35	0.0
η_{het}	0.16	2.14	0.11	-
$1/\gamma_{het}$	171.8	144.8	178.6	300.0
$\delta_{0,het}$	0.18	0.17	0.16	0.0
ρ_{hom}	3.47	4.55	1.55	0.0
ρ_{het}	3.5	2.18	3.26	0.0
d_{inf}	0.093 (17 μ g/week)	0.138 (25 μ g/week)	0.345 (62 μ g/week)	-

Table G.6: Estimated (mean posterior value) parameter vectors for patient #12, 18, and 32, when the assimilation time is $T = 600$. For patient #32, we present the results when considering a measurement of the CF at time $t = 508 > 300$. In addition, we present the estimated minimal IFN α dose d_{inf} (see § C.2.4). The last column gives an indication of the values we would have for WT cells, or more broadly, for cell populations which, even if mutated, would not have a proliferative advantage compared to WT cells, neither be targeted by IFN α .

G.2 Optimization

For the optimization part of this study, we consider having observed the patient (either patient #12, 18, or 32) before $T = 600$ days. We estimated their parameter vector from the observations before that time, as reported in Tab. G.6. In this section, we aim to optimize the therapy from time $T = 600$ days.

G.2.1 Patient #12

We consider $T = 600$ days. Using the observations of patient #12 before that time (7 VAF observations, two clonal architectures: the first at $t = 0$, the second at $t = 287$ days) and the prior knowledge obtained from the 18 remaining patients, we estimated the posterior distribution of the model parameters previously.

We now consider that the actual parameter vector is the one having been estimated (see Tab. G.6) and study how different dose strategies impact the response to the treatment after $T = 600$ days. With the estimated parameter vector, the minimal dose - under which the treatment would not target sufficiently the mutated HSCs, resulting in a relapse - is estimated to be equal to $d_{inf} = 0.093$, that is, 17 μ g/week. We consider different scenarios for how drug toxicity increases as a function of the dose, as presented in § 2.3.3 in the main text.

For each scenario, we study our three therapeutic strategies - presented in §2.3.2 in the main text - for which we find the parameters (e.g., the choice of the dose \bar{d}) that minimize the value of $M(\tau)$. This latter value corresponds to the toxicity-related amount of IFN α administered from T to the time τ , this latter being the time when the therapy could be interrupted (section D). We can compare the three therapeutic strategies for a given drug-toxicity relation and select the one with the best (i.e., the lowest) value for $M(\tau)$. Results are presented in Fig. G.9. It turns out that the constant strategy is found to be the most optimal for the linear and the concave scenario. For the convex scenario, the parameter vectors which minimize $M(\tau)$ for the periodic and the decreasing strategies are such that the three optimal strategies are almost the same. For the composite relation, $d_{inf} = 0.093 < d_{low} := 0.1$, d_{low} being the threshold below which we consider there is no toxicity (see E.4). Thus, it turns out in that case that each therapeutic strategy can be optimized to avoid administering a dose above the value of d_{low} , such that we end up with a zero value for $M(\tau)$. Under the hypothesis of the composite dose-toxicity relation, all optimal therapeutic strategies are close to the constant one. Thus, the constant strategy can also here be considered the most optimal.

For patient #12, the trade-off would be to treat them with a decreasing dose, starting at $T = 600$ days with $\bar{d} = 86\mu$ g/week, a decreasing factor $\lambda = 0.45$, and a period $L = 16$ months, until the treatment could be interrupted, here at age 77. This trade-off strategy would be:

- in the top 22 percent if the dose-toxicity relation were linear,
- in the top 46 percent if the dose-toxicity relation were convex,
- in the top 46 percent if the dose-toxicity relation were concave,
- and in the top 47 percent if the dose-toxicity relation were composite.

G.2.2 Patient #18

Using the observations of patient #18 before $T = 600$ (8 VAF observations, two clonal architectures: the first at $t = 0$, the second at $t = 248$ days) and the prior knowledge obtained from the 18 remaining patients, we previously estimated the posterior distribution of the model parameters.

We now consider that the actual parameter vector is the estimated one (see Tab. G.6) and study how different dose strategies impact the response of the treatment after $T = 600$ days. With the estimated parameter vector, the minimal dose is estimated to be equal to $d_{inf} = 0.138$, that is, $25\mu\text{g}/\text{week}$.

The constant strategy is found to be the most optimal for each of the four scenarios of dose-toxicity we study (Fig. G.10).

For patient #18, the trade-off would be to treat them with a constant dose of $61\mu\text{g}/\text{week}$, until the treatment could be interrupted, here at age 72. This trade-off strategy would be:

- in the top 0.09 percent if the dose-toxicity relation were linear,
- in the top 6.6 percent if the dose-toxicity relation were convex,
- in the top 7.5 percent if the dose-toxicity relation were concave,
- and in the top 3.9 percent if the dose-toxicity were composite.

G.2.3 Patient #32

The results related to patient #32 have been detailed in the main text (§3.2). Results are also displayed here in Fig. G.11.

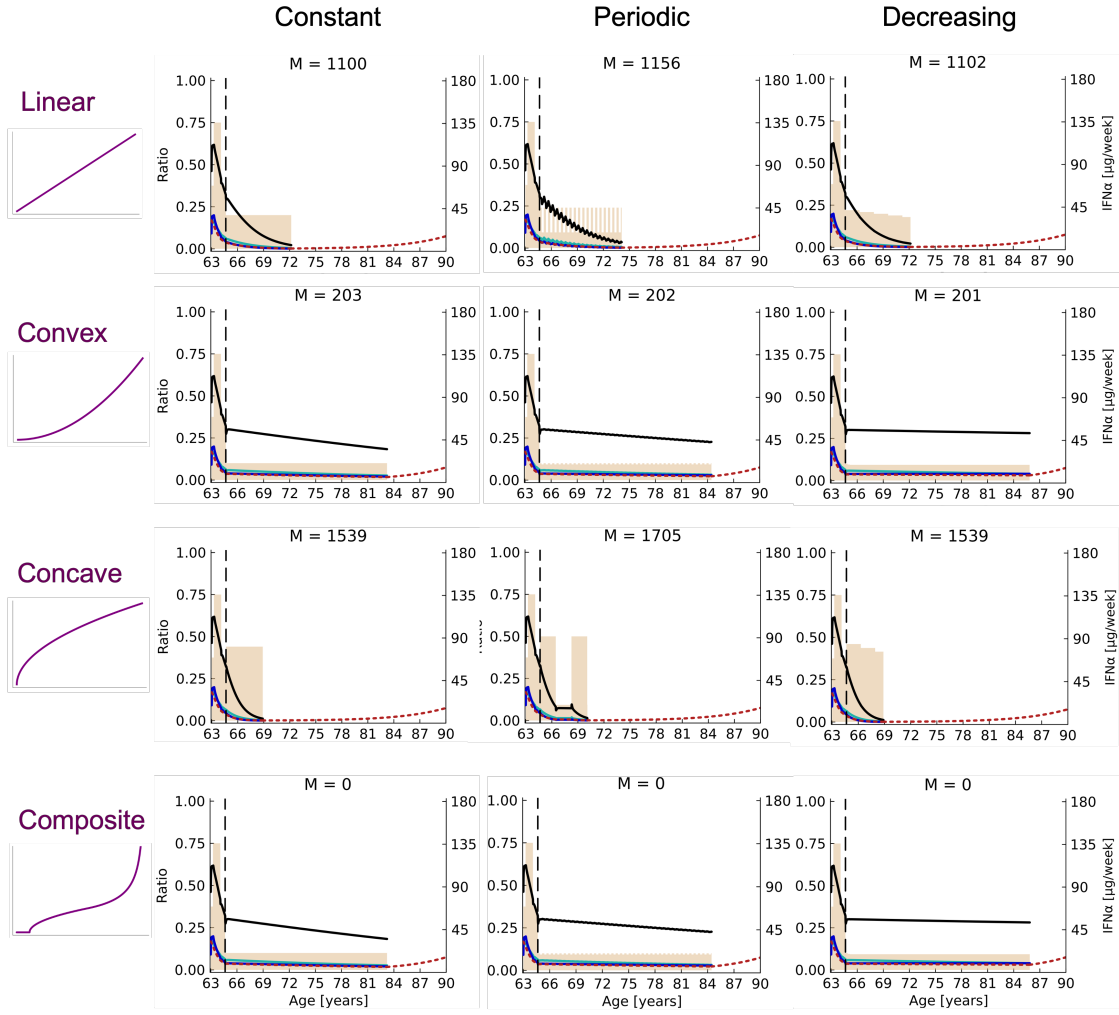


Figure G.9: Results of the treatment optimization for patient #12. Four hypothetical relationships between $\text{IFN}\alpha$ toxicity and the dose are studied: linear (at the top), convex (second line), concave (third line), and composite (at the bottom). For each of them, three therapeutic strategies - constant on the left, periodic on the middle, and decreasing on the right - are optimized to minimize $M(\tau)$ (denoted by M on the figures), which is the toxicity-related amount of $\text{IFN}\alpha$ administered from T (vertical dashed line) to the treatment interruption time τ (when the weekly dose of interferon - represented in beige - drops to zero). The dynamics of the VAF among mature cells (black line), the heterozygous CF among progenitors (green line), the homozygous CF among progenitors (blue line), and the VAF among HSCs (red dotted line) will differ for $t \geq T$ according to how the doses vary (depending on the considered therapeutic strategy and its parameter values). After the treatment interruption, only the dynamics of the VAF among HSCs is computed (and displayed), according to equation (D.3). All graphs are displayed until 90 years old, the age above which the VAF among HSCs should not exceed 7.5%, as explained in section D. M values are only comparable for a given dose-toxicity relation; the lower, the better.

If the actual dose-toxicity relation were linear, the constant strategy would be optimal for $\bar{d} = 36\mu\text{g}/\text{week}$ (with the associated age of treatment interruption $\tau = 72$ years and an optimal value of $M(\tau) = 1101$), the periodic strategy would be optimal for $\bar{d} = 43\mu\text{g}/\text{week}$ and $L = 3$ months, and the decreasing one for $\bar{d} = 40\mu\text{g}/\text{week}$, $L = 20$ months, and $\lambda = 0.95$. Among these three strategies, the constant one is the best.

If the actual dose-toxicity relation were convex, the constant strategy would be optimal for $\bar{d} = 18\mu\text{g}/\text{week}$, the periodic strategy would be optimal for $\bar{d} = 18\mu\text{g}/\text{week}$ and $L = 3$ months, and the decreasing one for $\bar{d} = 18\mu\text{g}/\text{week}$, $L = 3$ months, and $\lambda = 0.95$.

If the actual dose-toxicity relation were concave, the constant strategy would be optimal for $\bar{d} = 79\mu\text{g}/\text{week}$, the periodic strategy would be optimal for $\bar{d} = 90\mu\text{g}/\text{week}$ and $L = 22$ months, and the decreasing one for $\bar{d} = 83\mu\text{g}/\text{week}$, $L = 18$ months, and $\lambda = 0.95$. Among these three strategies, the constant and decreasing ones are both the best.

If the actual dose-toxicity were composite, with a low threshold $d_{low} = 0.1$ below which zero toxicity is assumed, then each strategy for which the dose is maintained below this threshold would lead to a zero value for $M(\tau)$.

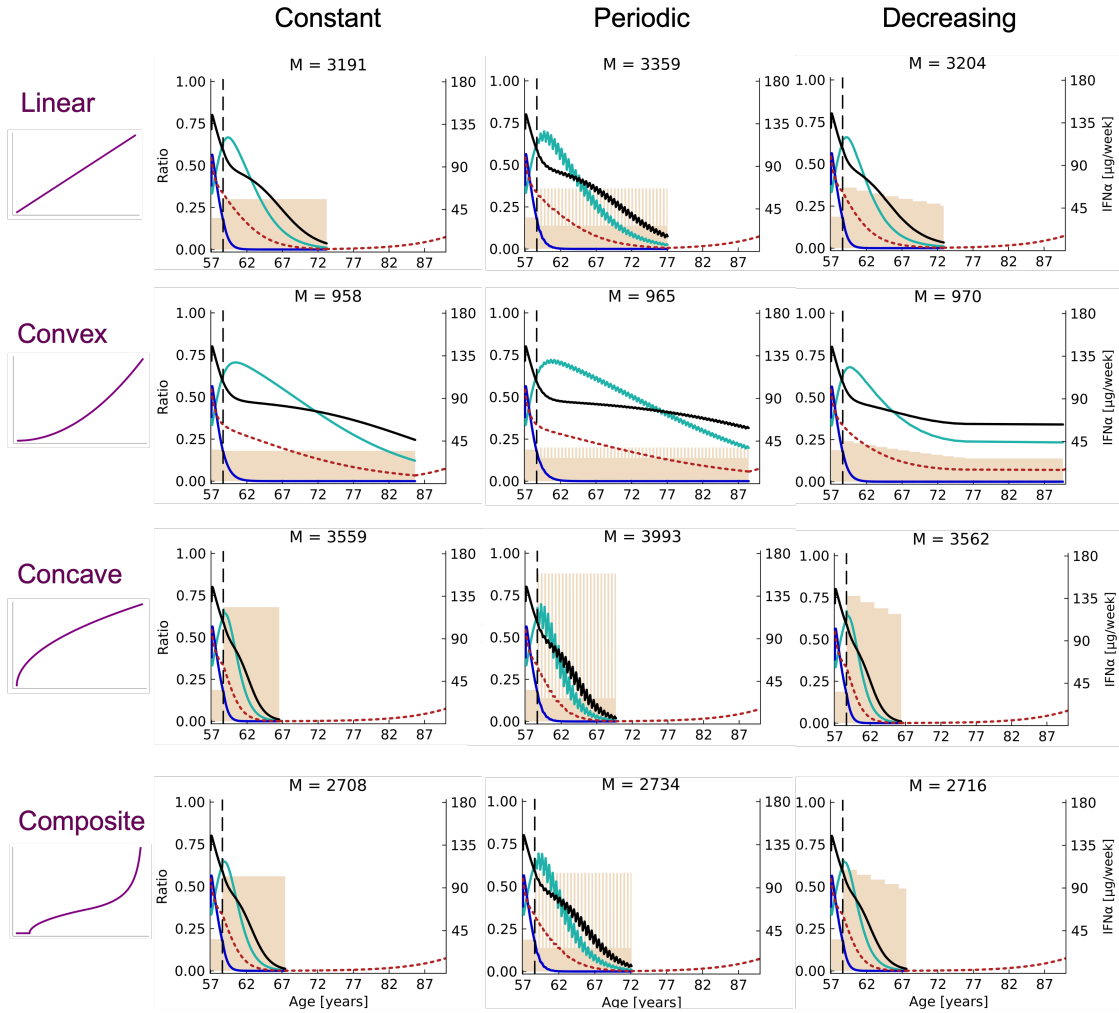


Figure G.10: Results of the treatment optimization for patient #18. Four hypothetical relationships between IFN α toxicity and the dose are studied: linear (at the top), convex (second line), concave (third line), and composite (at the bottom). For each of them, three therapeutic strategies - constant on the left, periodic on the middle, and decreasing on the right - are optimized to minimize $M(\tau)$ (denoted by M on the figures), which is the toxicity-related amount of IFN α administered from T (vertical dashed line) to the treatment interruption time τ (when the weekly dose of interferon - represented in beige - drops to zero). The dynamics of the VAF among mature cells (black line), the heterozygous CF among progenitors (green line), the homozygous CF among progenitors (blue line), and the VAF among HSCs (red dotted line) will differ for $t \geq T$ according to how the doses vary (depending on the considered therapeutic strategy and its parameter values). After the treatment interruption, only the dynamics of the VAF among HSCs is computed (and displayed), according to equation (D.3). All graphs are displayed until 90 years old, the age above which the VAF among HSCs should not exceed 7.5%, as explained in section D. M values are only comparable for a given dose-toxicity relation; the lower, the better.

If the actual dose-toxicity relation were linear, the constant strategy would be optimal for $\bar{d} = 54\mu\text{g}/\text{week}$ (with the associated age of treatment interruption $\tau = 73$ years and an optimal value of $M(\tau) = 3191$), the periodic strategy would be optimal for $\bar{d} = 65\mu\text{g}/\text{week}$ and $L = 3$ months, and the decreasing one for $\bar{d} = 65\mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best.

If the actual dose-toxicity relation were convex, the constant strategy would be optimal for $\bar{d} = 32\mu\text{g}/\text{week}$, the periodic strategy would be optimal for $\bar{d} = 36\mu\text{g}/\text{week}$ and $L = 3$ months, and the decreasing one for $\bar{d} = 43\mu\text{g}/\text{week}$, $L = 19$ months, and $\lambda = 0.95$.

If the actual dose-toxicity relation were concave, the constant strategy would be optimal for $\bar{d} = 122\mu\text{g}/\text{week}$, the periodic strategy would be optimal for $\bar{d} = 157\mu\text{g}/\text{week}$ and $L = 3$ months, and the decreasing one for $\bar{d} = 137\mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best.

If the actual dose-toxicity relation were composite, the constant strategy would be optimal for $\bar{d} = 100\mu\text{g}/\text{week}$, the periodic strategy would be optimal for $\bar{d} = 104\mu\text{g}/\text{week}$ and $L = 3$ months, and the decreasing one for $\bar{d} = 108\mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best.

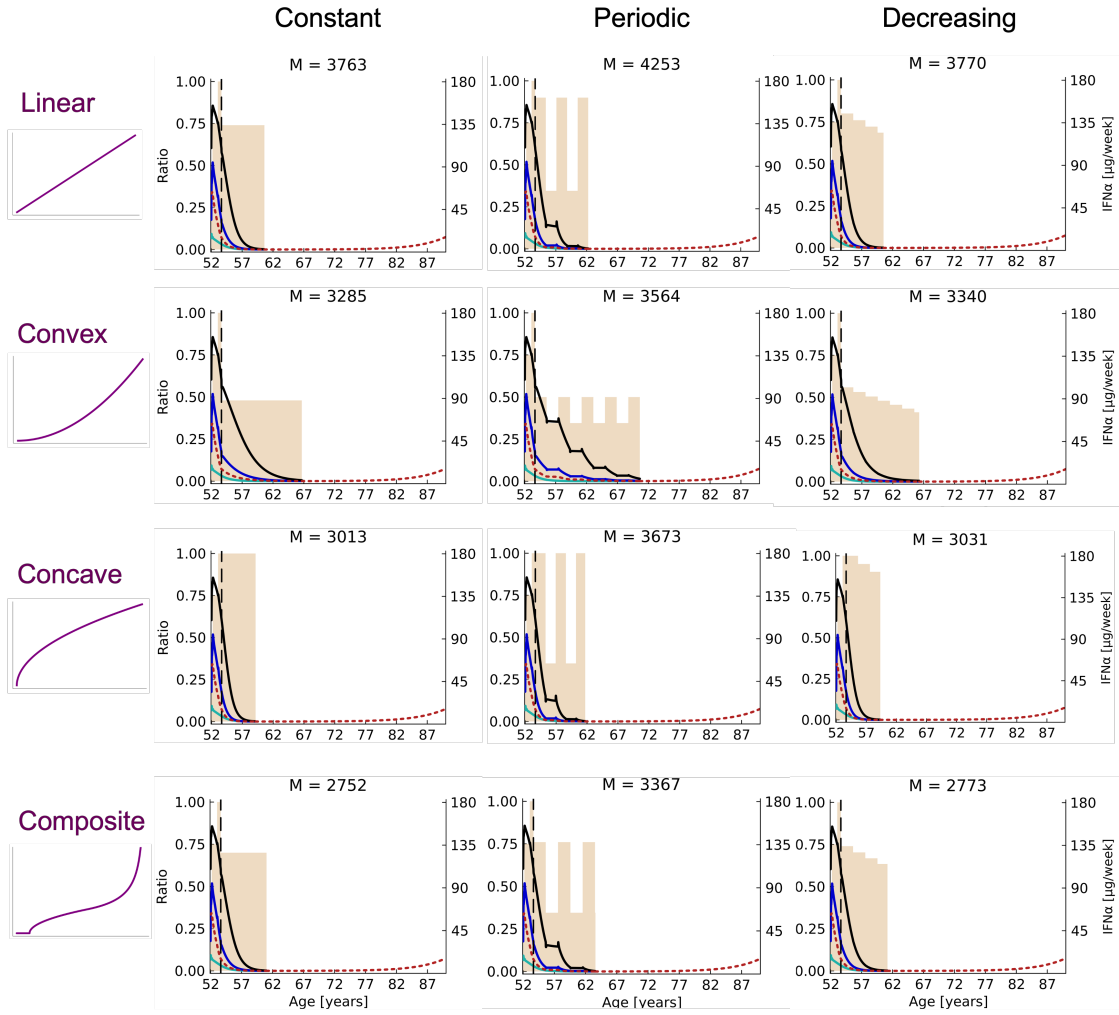


Figure G.11: Results of the treatment optimization for patient #32. Four hypothetical relationships between IFN α toxicity and the dose are studied: linear (at the top), convex (second line), concave (third line), and composite (at the bottom). For each of them, three therapeutic strategies - constant on the left, periodic on the middle, and decreasing on the right - are optimized to minimize $M(\tau)$ (denoted by M on the figures), which is the toxicity-related amount of IFN α administered from T (vertical dashed line) to the treatment interruption time τ (when the weekly dose of interferon - represented in beige - drops to zero). The dynamics of the VAF among mature cells (black line), the heterozygous CF among progenitors (green line), the homozygous CF among progenitors (blue line), and the VAF among HSCs (red dotted line) will differ for $t \geq T$ according to how the doses vary (depending on the considered therapeutic strategy and its parameter values). After the treatment interruption, only the dynamics of the VAF among HSCs is computed (and displayed), according to equation (D.3). All graphs are displayed until 90 years old, the age above which the VAF among HSCs should not exceed 7.5%, as explained in §D. M values are only comparable for a given dose-toxicity relation; the lower, the better.

If the actual dose-toxicity relation were linear, the constant strategy would be optimal for $\bar{d} = 135\mu\text{g}/\text{week}$ (with the associated age of treatment interruption $\tau = 61$ years and an optimal value of $M(\tau) = 3764$), the periodic strategy would be optimal for $\bar{d} = 153\mu\text{g}/\text{week}$ and $L = 22$ months, and the decreasing one for $\bar{d} = 144\mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best.

If the actual dose-toxicity relation were convex, the constant strategy would be optimal for $\bar{d} = 86\mu\text{g}/\text{week}$, the periodic strategy would be optimal for $\bar{d} = 90\mu\text{g}/\text{week}$ and $L = 23$ months, and the decreasing one for $\bar{d} = 100\mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best.

If the actual dose-toxicity relation were concave, the constant strategy would be optimal for $\bar{d} = 180\mu\text{g}/\text{week}$, the periodic strategy would be optimal for $\bar{d} = 180\mu\text{g}/\text{week}$ and $L = 20$ months, and the decreasing one for $\bar{d} = 180\mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best.

If the actual dose-toxicity relation were composite, the constant strategy would be optimal for $\bar{d} = 126\mu\text{g}/\text{week}$, the periodic strategy would be optimal for $\bar{d} = 137\mu\text{g}/\text{week}$ and $L = 2$ years, and the decreasing one for $\bar{d} = 133\mu\text{g}/\text{week}$, $L = 2$ years, and $\lambda = 0.95$. Among these three strategies, the constant one is the best.

G.3 Experimental Design

Our capacity to recommend good therapeutic strategies to clinicians depends on our capacity to correctly infer the model parameters of new patients, from only minimal observations. We showed in section G.1 that two observations of the clonal architecture and several VAF measurements before 600 days were sufficient to get accurate predictions. But we also illustrated with patient #32 that the quality of the predictions could highly depend on when the clonal architecture is measured (Section G.1.4). In this section, we aim to investigate whether the choice of the timing for measuring the heterozygous and homozygous CF could be rationalized. We apply the method we described in section F, and present first the results we get for patient #12, then the results for patients #18 and #32.

We consider $T = 600$ days. Using the clinical observations of patient #12 before that time (7 VAF observations, two clonal architectures: one at $t = 0$, the second one at $t = 287$) and the prior knowledge obtained from the 18 remaining patients, we estimated the posterior distribution of the model parameters and showed that we could make good predictions (section G.1.1).

We now consider that the actual parameter vector is the one having been estimated (mean of the posterior distribution). From it, we get the theoretical dynamics of the VAF and the heterozygous and homozygous CF over $[0, 600]$ days (solid lines on the top of Fig. G.12). We simulate 600 different synthetic datasets $\mathcal{D}_{T_{obs}}$. All of them have in common 7 VAF values (pseudo-observations) - obtained from the theoretical values by adding some noise - and the CF (pseudo-observation) at the initial time. All of them differ in the time point T_{obs} of the second CF (pseudo-observation). For each value of $T_{obs} \in \{1, 2, \dots, 600\}$, we estimate the parameter vector that maximizes the posterior density function, and infer the associated CF dynamics. We can then evaluate the error between the estimated and the theoretical CF dynamics, and find for which T_{obs} the error is minimal. Results for patient #12 are presented in Fig. G.12. For them, we find that the pseudo-observation time that would result in the most accurate parameter estimation is $T_{obs} = 100$ days, that is, about three months after the start of the therapy. This observation time corresponds to when the theoretical VAF and CF for patient #12 are maximal over $[0, 600]$, and also when there is a dose increase. However, when applying the same methods with patients #18 and #32, we find that it would be better to observe the clonal architecture later, at 476 days after the start of the therapy for patient #18 (Fig. G.13), when, for patient #32, the error between the theoretical and inferred CF is continuously decreasing when increasing the pseudo-observation time (Fig. G.14.) We are therefore unable to identify a common trend from the study of these three patients and only show that there would be better timings than others, which would be potentially patient-dependent.




Figure G.12: Optimal experimental design for patient #12. From the theoretical VAF dynamics (solid black line) of patient #12, we simulate noisy VAF pseudo-observations (black square). From the theoretical dynamics of the heterozygous (solid green line) and homozygous (solid blue line), we consider that we observe (pseudo-observations) the CF at the initial time and also at an additional observation time T_{obs} . In the bottom (purple line), we show the error between the theoretical and the inferred CF (L1-norm) according to the choice of T_{obs} . We see that the error would be minimal for $T_{obs} = 100$. On the top, the dashed lines correspond to the inferred dynamics when the synthetic dataset includes the CF pseudo-observation at $T_{obs} = 100$ days.




Figure G.13: Optimal experimental design for patient #18. From the theoretical VAF dynamics (solid black line) of patient #18, we simulate noisy VAF pseudo-observations (black square). From the theoretical dynamics of the heterozygous (solid green line) and homozygous (solid blue line), we consider that we observe the CF (pseudo-observations) at the initial time and also at an additional observation time T_{obs} . In the bottom (purple line), we show the error between the theoretical and the inferred CF (L1-norm) according to the choice of T_{obs} . We see that the error would be minimal for $T_{obs} = 476$. On the top, the dashed lines correspond to the inferred dynamics when the simulated dataset includes the CF pseudo-observation at $T_{obs} = 476$ days.

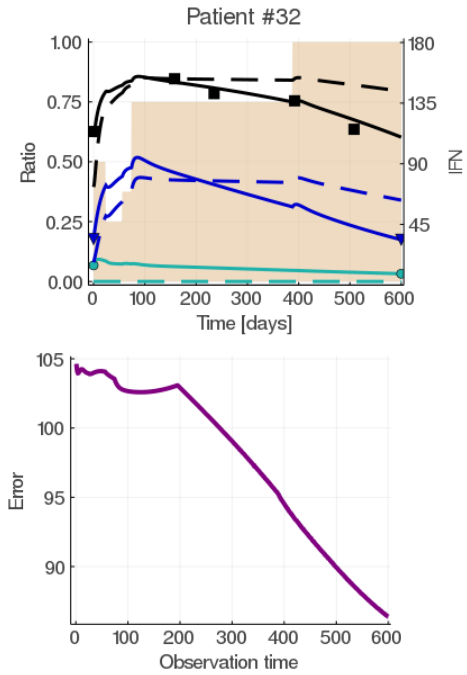


Figure G.14: Optimal experimental design for patient #32. From the theoretical VAF dynamics (solid black line) of patient #32, we simulate noisy VAF pseudo-observations (black square). From the theoretical dynamics of the heterozygous (solid green line) and homozygous (solid blue line), we consider that we observe (pseudo-observations) the CF at the initial time and also at an additional observation time T_{obs} . In the bottom (purple line), we show the error between the theoretical and the inferred CF (L1-norm) according to the choice of T_{obs} . We see that the error continuously decreases with T_{obs} . On the top, the dashed lines correspond to the inferred dynamics when the simulated dataset includes the CF pseudo-observation at $T_{obs} = 600$ days.

References

- [1] CATLIN, S. N., ABKOWITZ, J. L., AND GUTTORP, P. Statistical inference in a two-compartment model for hematopoiesis. *Biometrics* 57, 2 (2001), 546–553.
- [2] DUPONT, S., MASSÉ, A., JAMES, C., TEYSSANDIER, I., LÉCLUSE, Y., LARBRET, F., UGO, V., SAULNIER, P., KOSCIELNY, S., LE COUÉDIC, J. P., ET AL. The jak2 617v_i f mutation triggers erythropoietin hypersensitivity and terminal erythroid amplification in primary cells from patients with polycythemia vera. *Blood, The Journal of the American Society of Hematology* 110, 3 (2007), 1013–1021.
- [3] GEMAN, S., AND GEMAN, D. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, 6 (1984), 721–741.
- [4] HANSEN, N. The cma evolution strategy: a comparing review. *Towards a new evolutionary computation* (2006), 75–102.
- [5] HANSEN, N. The cma evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772* (2016).
- [6] HASAN, S., LACOUT, C., MARTY, C., CUINGNET, M., SOLARY, E., VAINCHENKER, W., AND VILLEVAL, J.-L. Jak2v617f expression in mice amplifies early hematopoietic cells and gives them a competitive advantage that is hampered by ifn α . *Blood, The Journal of the American Society of Hematology* 122, 8 (2013), 1464–1477.
- [7] HASTINGS, W. K. Monte carlo sampling methods using markov chains and their applications.
- [8] HERMANGE, G., RAKOTONIRAINY, A., BENTRIOU, M., TISSERAND, A., EL-KHOURY, M., GIRODON, F., MARZAC, C., VAINCHENKER, W., PLO, I., AND COURNÈDE, P.-H. Inferring the initiation and development of myeloproliferative neoplasms. *Proceedings of the National Academy of Sciences* 119, 37 (2022), e2120374119.
- [9] HERMANGE, G., VAINCHENKER, W., PLO, I., AND COURNÈDE, P.-H. Mathematical modelling, selection and hierarchical inference to determine the minimal dose in ifn α therapy against myeloproliferative neoplasms. *arXiv preprint arXiv:2112.10688* (2021).
- [10] KILADJIAN, J.-J. Final results of ruxopeg, a phase 1/2 adaptive randomized trial of ruxolitinib (rux) and pegylated interferon alpha (ifna) 2a in patients with myelofibrosis (mf). In *64th ASH Annual Meeting and Exposition* (2022), ASH.
- [11] KONTIS, V., BENNETT, J. E., MATHERS, C. D., LI, G., FOREMAN, K., AND EZZATI, M. Future life expectancy in 35 industrialised countries: projections with a bayesian model ensemble. *The Lancet* 389, 10076 (2017), 1323–1335.
- [12] LEE-SIX, H., ØBRO, N. F., SHEPHERD, M. S., GROSSMANN, S., DAWSON, K., BELMONTE, M., OSBORNE, R. J., HUNTLY, B. J., MARTINCORENA, I., ANDERSON, E., ET AL. Population dynamics of normal human blood inferred from somatic mutations. *Nature* 561, 7724 (2018), 473–478.
- [13] MANN, H. B., AND WHITNEY, D. R. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics* (1947), 50–60.
- [14] MICHOR, F., HUGHES, T. P., IWASA, Y., BRANFORD, S., SHAH, N. P., SAWYERS, C. L., AND NOWAK, M. A. Dynamics of chronic myeloid leukaemia. *Nature* 435, 7046 (2005), 1267–1270.
- [15] MOSCA, M., HERMANGE, G., TISSERAND, A., NOBLE, R., MARZAC, C., MARTY, C., LE SUEUR, C., CAMPARIO, H., VERTENOEIL, G., EL-KHOURY, M., ET AL. Inferring the dynamics of mutated hematopoietic stem and progenitor cells induced by ifn α in myeloproliferative neoplasms. *Blood* (2021).
- [16] VAN EGEREN, D., ESCABI, J., NGUYEN, M., LIU, S., REILLY, C. R., PATEL, S., KAMAZ, B., KALYVA, M., DEANGELO, D. J., GALINSKY, I., ET AL. Reconstructing the lineage histories and differentiation trajectories of individual cancer cells in myeloproliferative neoplasms. *Cell stem cell* 28, 3 (2021), 514–523.

- [17] WILLIAMS, N., LEE, J., MITCHELL, E., MOORE, L., BAXTER, E. J., HEWINSON, J., DAWSON, K. J., MENZIES, A., GODFREY, A. L., GREEN, A. R., ET AL. Life histories of myeloproliferative neoplasms inferred from phylogenies. *Nature* 602, 7895 (2022), 162–168.