Cancer-immune coevolution dictated by antigenic mutation accumulation

Long Wang[1†], Christo Morison[2†], and Weini Huang[1,2*]

[1]Group of Theoretical Biology, Innovation Center for Evolutionary Synthetic Biology, School of Life Sciences, Sun Yat-sen University, Guangzhou, China
[2]School of Mathematical Sciences, Queen Mary University of London, UK
[†]These authors contributed equally to this work
[*]Corresponding author: weini.huang@qmul.ac.uk

9th September 2024

Abstract

The immune system is one of the first lines of defence against the emergence of cancer. When effector cells attempt to suppress the tumour, the cancer cells can respond in kind by evolving methods of escape or inhibition. Knowledge of this coevolutionary system and the selection taking place within it can help us understand tumour-immune dynamics both during tumorigenesis but also when treatments such as immunotherapies are applied. Here, we present an individual-based branching process model of mutation accumulation, where random mutations arising in cancer cells trigger corresponding specialised immune responses. Different from previous research, we explicitly model interactions between cancer and effector cells, while incorporating stochastic effects, which are especially important for the expansion and extinction of small populations. We find that the parameters governing interactions between the cancer and effector cells induce different outcomes of tumour progress, such as suppression and evasion. While it is hard to measure the cancer-immune dynamics directly in patients, genetic information of the cancer may indicate the presence of such interactions. Our model demonstrates signatures of selection in sequencing-derived summary statistics, such as the single-cell mutational burden. Thus, bulk and single-cell sequencing of a tumour may give information about the coevolutionary dynamics.

Keywords: cancer-immune coevolution, mutation accumulation, stochastic modelling, site frequency spectrum, single-cell mutational burden distribution

Introduction

The immune system does not only protect our body from infectious diseases caused by various pathogens, but also gives the first response against threats emerging within the body, such as cancer. Immune cells try to identify and eliminate tumour cells, which may express antigens not found on normal cells. Meanwhile, tumour cells attempt to hide or evade from immune surveillance [1]. From an ecological perspective, this antagonistic relationship gives rise to complex dynamics between a tumour and its microenvironment [2, 3]. On an evolutionary level, these two cells types, though belonging to the same organism, coevolve, each adapting to genotypic and phenotypic changes in the other population [4]. We are interested in the eco-evolutionary dynamics of genotype-specific interactions between cancer and immune cells arising from a continuous introduction of new antigens with the context of immune surveillance and escape.

Dunn *et al.* described the battle between the immune system and emergent tumours in three stages, termed the three Es of cancer immunoediting [5]. The first, Elimination, was formulated by Burnet as the immunosurveillance hypothesis, which states that the immune system can win this battle and eliminate small cancers [6, 7]. According to the "bad luck" hypothesis [8], this is a frequent occurrence: only stochastically does the immune system allow cancers to sneak through [4]. Should the cancer do so, it enters Equilibrium. Indeed, some cancers take years to grow to detectable size [5], and there is evidence for small persistent tumours coevolving with the immune system [9]. Finally, Escape, when the cancer evolves mechanisms of evasion, and grows to a size that can be detected [10]. Deciphering genetic footprints from this coevolutionary process is of utmost clinical relevance from prognosis to treatment [11], especially since current measures such as immune infiltration, evidence of immune escape and tumour mutational burden (TMB) are not foolproof markers of immunotherapeutic success and overall survival [12].

Mutations accumulated during cancer evolution increase intratumour heterogeneity, providing a wide landscape of genotypes to improve the tumour's persistence [13]. Propensities for mutating more rapidly, increasing growth rate, developing metastases, and acquiring resistance to eventual treatment are possible consequences of this accrual of mutations [1]. However, mutations may also alert the immune system of the presence of tumour cells and initiate a suppressive response. This occurs when a mutation acquired by a cancer cell (called an antigenic mutation) results in the presentation by human leukocyte antigen (HLA) of immunogenic peptides called neoantigens at the cell surface [14]. These neoantigens are recognisable by cytotoxic T lymphocytes (CTLs), specialised effector cells of the adaptive immune system, which bind to the presented neoantigen via a T cell receptor at their surface and kill the targeted cancer cell [15]. While antigenic mutations are neutral in the absence of an immune response [16], their fitness in general depends on the likelihood of neoantigen presentation and recognition by effector cells [17].

Antigenic mutations so targeted by the immune system are thus under negative selective pressure [18]. The killing of cancer cells carrying antigenic mutations releases further neoantigens, leading to a cyclic process of generating immunity to cancer [19]. Cancer cells, however, may in turn combat the immune response via several mechanisms. Cancer cells can inhibit effector cells, such as by expression of programmed cell death-ligand 1, which normally presents in healthy cells to stop the attack of immune cells and can exhaust CTLs interacting with cancer cells that carry it [20, 21]. Some cancer cells may escape the immune response, such as by reducing neoantigen presentation through down-regulating HLA [14] or by immune-editing and losing antigenic mutations due to the immune negative selection [22]. Cancer cells can even develop immune exclusion, physically restricting the immune cells' access to the tumour [1]. These processes may arise individually or in concert, and they mean the immune system plays a crucial role in shaping a tumour's evolution [1, 10].

Immunotherapies leverage the immune system to reverse the evolution of evasion by mitigating or interfering with these processes [21]. Cytokine-based and tumour-infiltrating leukocyte-based immunotherapies help increase the effector cell population [15, 23]; immune checkpoint blockade therapies restore the effector response to immune-escaped cancer cells [12, 19]; other immunotherapies simply boost the ability of CTLs to kill tumour cells [15, 24]. Understanding of the coevolution between cancer and immune cells, along with the cell-to-cell interactions that drive it, may inform when therapies will succeed—and why.

There has been a long history of studying antagonistic coevolution experimentally and theoretically [25–28], with extensive literature on mathematical models of tumour-immune interactions [21, 29, 30]. Deterministic models often describe the antagonistic relationship between cancer and immune cells as obeying Lotka-Volterra dynamics, with the immune system predating on its tumour prey [31, 32]. Stochastic modelling of the tumour-immune system has been explored by *e.g.* George and Levine, who characterised escape as a random process with sequential mutations [4, 33] and subsequently framed cancer evolution as an active optimisation process in response to an evolving immune landscape [34, 35]. In concert with patient sequencing data, Lakatos and colleagues proposed and applied a model of random antigenic mutation accumulation in order to describe the negative selection that neoantigens undergo, resulting in neutral-like evolutionary dynamics [10, 11]. Recently, Chen & Xie *et al.* included negative frequency dependence in a similar model of antigenic mutation accumulation, so that tumours are only immunogenic when a sufficiently-large proportion of their cells present neoantigens, predicting that tumours undergoing this frequency-dependent selection have poorer treatment outcomes than their negative selection counterparts [36].

Coevolution between the immune system and the threats it faces has often been studied in a gene-for-gene framework, which centres on the genetic makeup of multiple populations being tracked [37]. However, there is a dearth of stochastic models that explore the explicit evolutionary dynamics of both tumour and immune cell populations [34]: the aforementioned models either encapsulate the immune response into a selection parameter s, implicitly assuming that effector cells react perfectly and instantaneously to a cancerous threat [11, 36]; or, they are deterministic and thus miss out the critical impact of random processes on small population sizes, while omitting direct genetic information with which sequencing data can be compared [38].

We address this gap in the literature by presenting and analysing a novel stochastic coevolutionary model of tumour-immune dynamics. As in Lakatos *et al.*'s model, antigenic mutations accrue in cancer cells undergoing a branching process and are negatively selected against by the immune system [11]. The adaptive immune system is represented by specialised effector populations that react to emergent neoantigens [39], leading to complex dynamics on both tumour and immune fronts. We focus on interactions between effector cells and their targets, incorporating both active and passive recruitment of effector cells, killing of cancer cells, and inhibition of effectors by cancer cells, while including explicit mechanisms of escape. In particular, we inspect how these interactions are central in determining the evolution of the system. Stochastic simulations allow us to characterise the various dynamics and outcomes that emerge, along with informing the impact of immunotherapy.

Model

Birth-death process with mutation accumulation and escape

Consider a population of cancer cells undergoing a stochastic branching process with per-capita birth rate b and per-capita death rate d (see Figure 1A). We assume births and deaths happen randomly based on their rates, with exponentially-distributed wait times between events, as in the Gillespie algorithm [40]. At each cancer cell division, daughter cells inherit their mother cell's mutations and acquire a random number of new mutations, drawn from a Poisson distribution with mean λ [41], as shown in Figure 1B. Similar to Lakatos *et al.*, we focus on the exomic region [11], where the value of λ ranges between 1 and 10. All mutations are unique, as in the infinite sites approximation, where the exome is considered long enough for two co-occurring point mutations to be negligible [42]. We distinguish between two types of mutations: with probability p_a, a mutation is antigenic and can be recognised by given effector cells, and with probability $1 - p_a$ it is neutral [11, 15]. During a division, there is a probability p_e that an antigenic mutation escapes the immune system, therefore making the cell possessing it (and all of its descendants [36, 43]) undetectable to the immune system (see Figure 1B).

Figure 1. Cancer-immune stochastic model. **A** Cancer cells (red) stochastically divide and die with rates b and d, respectively (here $\varnothing$ represents no cell). **B** During a division event, daughter cells inherit all of their mother cell's antigenic and neutral mutations, depicted by numbers (where underlined numbers are antigenic). Cells carrying antigenic mutations have probability p_e to escape and become neutral, as shown in the lower daughter cell with mutation 2. Each daughter cell also acquires a random number (drawn from a Poisson distribution with mean λ) of new mutations, where each mutation is antigenic with probability p_a and neutral with probability $1 - p_a$. **C** For each antigenic mutation i present in the system, a corresponding effector cell population E_i exists (blue), which grows with constant rate B and shrinks with per-capita death rate D. **D** Antigenic mutations in cancer cells (such as i and k) display unique neoantigens at the cell surface, whereas neutral mutations (such as j) do not. The neoantigens can be identified by specialised effector cells, which can only interact with the corresponding cancer cells. **E** When a cancer cell carrying the antigenic mutation i meets an effector cell of type i, three outcomes are possible: active recruitment of another effector cell of type i with rate α_i, killing of the cancer cell with rate β_i and inhibition/exhaustion of the effector cell with rate γ_i. The expressions for the rates are found in equation (1).

Specialised effector response and cancer-immune interactions

Neoantigens trigger responses from the adaptive immune system: each antigenic mutation i calls a unique, specialised effector population E_i, as in gene-for-gene coevolution [37]. Whenever mutation i exists within the population, effector cells of type i are passively recruited from the body at constant rate B ($\varnothing \to (E_i)$, for $\varnothing$ the absence of a cell and (E_i) an effector cell of type i, with rate B), and die with per-capita rate D [44, 45], as shown in Figure 1C.

To each antigenic mutation i, we associate two random numbers, each drawn independently from an exponential distribution with mean 1: an antigenicity A_i, describing the propensity for an effector cell (E_i) to kill a cancer cell possessing mutation i, and an immunogenicity I_i, which relates to the rate at which active recruitment of more effector cells of type i occurs during the specified tumour-immune interaction (see equation (1) and Figure 1E). These can be thought to encapsulate the probability of an antigenic mutation leading to the presentation of neoantigens by HLA molecules and subsequent recognition by CTLs [17].

Correspondingly, when a cancer cell with an antigenic mutation presents a neoantigen at its surface, only effector cells of the corresponding population can interact with it [21], as shown in Figure 1D. Interactions between cancer cells possessing an antigenic mutation i and corresponding effector cells of type i have three possible outcomes: active recruitment with rate α_i, killing with rate β_i and inhibition/exhaustion with rate γ_i (see Figure 1E):

$$(C_i), (E_i) \to \begin{cases} (C_i), (E_i), (E_i) \\ (E_i) \\ (C_i) \end{cases} \quad \text{with rates} \quad \begin{aligned} \alpha_i &= \frac{\alpha_0 I_i}{1 + \alpha_0 h_\alpha C_i I_i} \\ \beta_i &= \beta_0 A_i \\ \gamma_i &= \gamma_0, \end{aligned} \tag{1}$$

where α_0, β_0, γ_0 and h_α are constants for the entire tumour. Note that the reactions of active recruitment and inhibition/exhaustion have opposite effects on the effector cell population; here, we assume $\alpha_0 > \gamma_0$, the opposite of which is rarely considered [46], as it leads to net decreasing effector populations upon interactions with cancer cells. However, active recruitment α_i obeys a type-II functional response rather than a type-I (that is, linear) response as inhibition/exhaustion γ_i [47, 48]. This implies that $(\alpha_i - \gamma_i)C_i$ is not linear in the cancer cell population C_i, whereas if α_i too satisfied a type-I response it would be. Instead, as shown in Figure A1 of the Appendix, $(\alpha_i - \gamma_i)C_i$ is positive for small populations C_i and negative once the clone C_i is large enough.

Mutational distributions

We are interested in genetic information relating to each of the two populations, which we will use to identify and measure the coevolution between effector and cancer cells. For the cancer cells, we define the following summary statistics: the site frequency S_j, denoting the number of antigenic mutations present in j cells, and the single-cell mutational burden B_k, the number of cells that possess k antigenic mutations. The conglomeration of these (for non-negative integers j and k, respectively) form the site frequency spectrum (SFS) and the single-cell mutational burden distribution (MBD). These satisfy $\sum_{j=1}^C jS_j = \sum_{k=1}^M kB_k$, where C denotes the total population of cancer cells and M denotes the total number of antigenic mutations [49]. We will write U for this quantity (that is, for either side of the previous equality), the total number of antigenic mutational occurrences.

For the effector population, we define the average antigenicity $\langle A \rangle$ and the average immunogenicity $\langle I \rangle$ as follows:

$$\langle A \rangle = \frac{\sum_{i=1}^M A_i E_i}{\sum_{i=1}^M E_i} \quad \text{and} \quad \langle I \rangle = \frac{\sum_{i=1}^M I_i E_i}{\sum_{i=1}^M E_i}. \tag{2}$$

Models that focus only on the cancer cell populations ascribe an antigenicity to the tumour itself [11, 36]; here, by considering both cancer and effector populations, we can also average the antigenicities and immunogenicities over all effector cells. By choosing to average over the effector cells as in equation (2), $\langle A \rangle$ and $\langle I \rangle$ become measures of the gene-for-gene immune response and its potency in fighting its tumour target, providing an additional angle on the coevolutionary system not present in single-population models.

We implement our model using a standard Gillespie algorithm [40]. A realisation ends when either the cancer cell population exceeds a threshold size of K (which we call *no suppression* and write T_K for the time that this occurs); or, the cancer cell remains below K (which we call *suppression* and write T_{end} as the end of the simulation time). Because the model is stochastic, different outcomes (suppression or not) may arise from a same set of parameters and initial conditions.

Symbol	Baseline	Description
b, d	1, 0.1	Birth and death rates of a cancer cell
λ	1, 10	Mean number of exomic mutations acquired per division per daughter cell
p_a	0.075	Probability of a mutation being antigenic (rather than neutral)
p_e	10^{-6}	Probability of escape of an antigenic cancer cell into neutrality
C_i		Population of cancer cells carrying antigenic mutation i, with total cancer cell population C
E_i		Population of effector cells of type i, with total effector cell population E
B, D	0.2, 0.1	Passive recruitment and per-capita death of an effector cell
$\alpha_i, \beta_i, \gamma_i$	—, —, 10^{-3}	Active recruitment, killing and inhibition/exhaustion rates, respectively, for interactions between cancer and effector cells of type i
A_i, I_i		Antigenicity and immunogeniticy, respectively, of antigenic mutation i (each drawn from Exp(1)), with averages over the effector population $\langle A \rangle$ and $\langle I \rangle$
S_j		Site frequency: number of antigenic mutations occurring j times
B_k		Single-cell mutational burden: number of cancer cells with k antigenic mutations
M		Total number of antigenic mutations
U		Total antigenic mutational occurrences: $U = \sum_{j=1}^{C} jS_j = \sum_{k=1}^{M} kB_k$
K	3×10^4	Ending cancer population size for stochastic simulations
T_{end}	16	Maximum time for a simulation; if the cancer cell population never grows to size K, then the outcome is *suppression*, otherwise it is *no suppression* (at time T_K)

Table 1. Notation and baseline parameter values, which are used for all simulations unless otherwise specified.

Results

Cyclic dynamics between antigenicity and immunogenicity

Antagonism between effector cells and their cancer targets result in a range of complex dynamics. For elevated immune efficacy such as high values of α_0 and β_0 (the active recruitment and killing parameters), emerging tumours are suppressed in early stages. However, when the probability of escape p_e is large enough, then neutral-like dynamics ensue, as most cancer cells evade the immune response. While it is expected that negative selection of the immune system on new antigens will selectively prune cancer cells with more antigenic mutations [11], this process depends on the effector populations themselves, through the interactions formulated in equation (1). Since here we explicitly model both tumour and immune cell types, we are able to explore their population dynamics in concert.

Indeed, we can clearly observe a dynamical response between the cancer and effector cells in single evolutionary trajectories. Figure 2 depicts two representative realisations of our stochastic simulations, one for a lower mutation rate ($\lambda = 1$) and one for a high mutation rate ($\lambda = 10$), the latter of which can be thought of as representing hyper-mutated tumours [11]. The population dynamics of Figures 2A and 2C show the effector population undergoing irregular spikes, out of phase with the cancer population, as seen in other coevolutionary models [50, 51]. Because of the specialised nature of our model, these arise when a single antigenic mutation i causes the rapid active recruitment of the corresponding effector cell population E_i to combat the subclone possessing mutation i during that time period. Once mutation i is eradicated from the cancer cell population, the effector cells E_i specialised to their neoantigen are removed from the system, as they no longer play a dynamical role and die out exponentially (see Figure 1C). The expected frequency and amplitude of these effector spikes can be approximated, as described in Section A8 of the Appendix.

We validate this by inspecting the coevolution of antigenicity and immunogenicity in the corresponding single realisations. In Figures 2B and 2D, we can see that the effector spikes arise due to the emergence of a mutation i that has a large immunogenicity I_i (significantly greater than the mean $\langle I \rangle$), as demonstrated in equation (A18) of the Appendix. Unexpectedly, in addition to this elevated immunogenicity, Figures 2B and 2D imply that these spikes also arise for mutations with low antigenicity $A_i < \langle A \rangle$. These fluctuations in $\langle I \rangle$ and $\langle A \rangle$ are coupled to the spikes in the population size; had a mutation emerged with a large antigenicity, it would have been quickly eradicated, and so the effector population size would not have grown sufficiently to be identified as a spike. As expected, the higher mutation rate ($\lambda = 10$) leads to more cycles of spikes in the cancer-effector dynamics as well as in the coevolution of antigenicity and immunogenicity dynamics.

Figure 2. Single stochastic realisations for $\lambda = 1$ (**A-B**) and $\lambda = 10$ (**C-D**). **A** & **C** Population dynamics, where red and blue lines depict total cancer and effector cell populations. **B** & **D** Average immunogenicity (yellow line) and antigenicity (orange line) in the effector population for the corresponding realisation. Interaction parameters: $\alpha_0 = 0.03$ and $\beta_0 = 0.3$ (**A-B**) and $\alpha_0 = 0.005$ and $\beta_0 = 0.01$ (**C-D**).

Interactions dictate outcome of tumour progress

The outcome of the coevolution between the immune system and a tumour depends strongly on the interaction parameters between effector and cancer cells (equation (1)). Figures 3A and 3D illustrate heat maps for the tumour suppression proportion across different values of the active effector recruitment rate α_0 and the cancer killing rate β_0. Due to the stochastic nature of the model, the outcome for a given parameter set is probabilistic. As expected, for higher values of α_0 and β_0, more effector cells are recruited and more cancer cells are killed, thus the suppression increased. While this pattern holds for different mutation rates (Figure 3), it is more dominating when mutation rate is high. For a higher mutation rate ($\lambda = 10$), the immune system is more effective since there are more antigenic mutations to target, resulting in a high suppression (see Figure 3D). For a lower mutation rate ($\lambda = 1$), however, only a small proportion of tumours are suppressed even with strong immune response (Figure 3A). This is because only a fraction of the population is antigenic (and thus undergoing immunoselection): observe Figures A6A-C versus A6D-F of the Appendix. It also worth noting that in this case most suppressed tumours go extinct at early times (Figure A5A) due to stochasticity.

When the active effector recruitment rate is smaller than the effector inhibiting rate $\alpha_0 < \gamma_0$ (Figure 3 red dashed line), the net outcome for effectors during an interaction is negative (though there are subtleties here due to the difference in functional responses; see Section A4 of the Appendix for details). This implies that the resulting suppressed tumours in this parameter region were either defeated by a passively-recruited effector population, or by effector types i corresponding to antigenic mutations with particularly high immunogenicities I_i. This is possible since I_i is drawn from an exponential distribution (and thus can be high), meaning that even when α_0 is low, the active recruitment can be greater than the inhibition/exhaustion by chance. This is exactly what we observed in our stochastic simulations. Under high mutation rate ($\lambda = 10$), most tumours are not suppressed under low interaction parameters α_0 and β_0 (Figure 3D), whereas for high interaction parameters most tumours go extinct (Figure 3E) rather than maintaining a slow growing pace (Figure 3F). Again, we see a similar pattern under a low mutation rate ($\lambda = 1$), though more weakly and with more noise.

Interestingly, for a high mutation rate there exists an intermediate range of interaction parameters that allows for tumours to neither go extinct nor to grow to capacity (Figure 3F). One further interpretation of Figure 3D is the presence of a killing threshold (here, at $\beta_0 \approx 10^{-1}$) above which tumours are suppressed, no matter the active recruitment rate α_0. The rest of the domain of Figure 3D then exhibits a much stronger dependence on the active recruitment rate α_0, in line with the results of Wilkie & Hahnfeldt [16].

Figure 3. Outcome heat maps for $\lambda = 1$ (**A**-**C**) and $\lambda = 10$ (**D**-**F**) for tumours interacting with an immune system characterised by cancer cell-effector cell interaction parameters α_0 (active recruitment of effectors) and β_0 (killing of cancer cells). **A** & **D** Proportion of suppressed tumours. **B** & **E** Proportion of extinct tumours. **C** & **F** Proportion of slow-growing tumours: tumours that are suppressed but do not go extinct. Points a and b (c and d for $\lambda = 10$) label parameter sets of low and high immune effectiveness, respectively. The red dashed line denotes γ_0 (here $\gamma_0 = 10^{-3}$), the rate of inhibition/exhaustion (see Section A4 for further discussion); see the Discussion section for details of point e. All parameter values not specified here are listed in Table 1.

Genetic markers of selection

Stochastic mutation accumulation in an exponentially growing population has been widely studied [52, 53]. When the immune system has little impact on the cancer cell population, therefore, it is unsurprising to see consistent increases in the average number of mutations per cancer cell, as in Figure 4A. The theoretical expectation of neutrally accumulated mutations can be approximated by the average flux of number of new mutations entering the system, $2bt$ [54], as plotted. (Note that a correction for large times has recently been shown to apply [55], which helps explain why the dashed line over-estimates the simulated data in Figure 4.) When the effector population is selectively killing, however, possessing more antigenic mutations makes a cancer cell more likely to get killed, so the average number of antigenic mutations per cancer cell does not continuously increase with the population, as shown in Figure 4B. This effect of selection is even more pronounced in the case of a higher mutation rate ($\lambda = 10$), as shown in Figure A7 of the Appendix.

Figure 4. Average number of antigenic mutations per cell (solid oranges lines) for several representative realisations when $\lambda = 1$. Theoretical prediction for the accumulation of neutral mutations per cell in an exponentially-growing population shown in grey dashed line. **A** Low immune effect: $\alpha_0 = 0.002$ and $\beta_0 = 0.001$. **B** High immune effect: $\alpha_0 = 0.03$ and $\beta_0 = 0.3$. (Parameter sets chosen as points a and b from Figure 3A.)

The average number of antigenic mutations per cell is the mean of the MBD: $\langle B \rangle = U/C$, which, like the SFS, can be extracted from bulk data. However, single-cell sequencing data can also provide information on the MBD overall, which can be used in combination with bulk information to infer what selection is taking place in the system [56]. As discussed in Section A5 of the Appendix, the expected neutral distributions of the SFS and the MBD have been solved [49, 57], so divergence from these may demonstrate the presence of selection and the strength of cancer-immune interactions. In Figure 5, the SFS and MBD averaged over 100 realisations are plotted in conjunction with the corresponding theoretical predictions (black dashed line) for the case with no immune response (that is, where all mutations are neutral). The first and third rows (green points) represent neutral mutations, whereas the second and fourth rows (orange points) measure antigenic mutations; for the MBDs, the mean (that is, the value U/C) was plotted in dashed vertical lines for both the simulated data (in green and orange) and the theoretical predictions (in grey), which were calculated with equations (A12) and (A13) of the Appendix.

We notice that in the case of low immune effectiveness (Figures 5A-D), there is little deviation from the neutral expectation. When the immune system plays a larger role, however, the distinction is significant, as in Figures 5E-H. In particular, the cells with more antigenic mutations were more selectively killed by the immune system, so that the tail of the MBD is smaller and the mean is shifted down, as seen most prominently in Figure 5H. As before, when the mutation rate is higher ($\lambda = 10$), these effects are more striking, as shown in Figure A8 of the Appendix. On the other hand, the SFS shows limited difference from its neutral expectation (Figures 5A and 5E), reiterating the importance of integrating single-cell data with bulk sequencing data in identifying immune effects. Under stronger negative selection, we expect a depletion of antigenic mutations from the population and thus lower site frequencies than the theoretical prediction [11]. This is visible as a slight depression of the data compared to the neutral prediction in Figure 5F, though such an observation is much clearer for a faster immune response, as discussed in Section A9.

Figure 5. Genetic markers of selection and cancer-immune interactions for $\lambda = 1$. We present the SFS of neutral (**A-B**) and antigenic mutations (**E-F**) and the MBD of neutral (**C-D**) and antigenic mutations (**G-H**). Left-column panels depict low immune effectiveness ($\alpha_0 = 0.002$ and $\beta_0 = 0.001$) and right-column panels depict high immune effectiveness ($\alpha_0 = 0.03$ and $\beta_0 = 0.3$). Black dashed lines are the theoretical predictions in the absence of an immune response. The dashed vertical lines represent the means of the MBDs in panels **C**, **D**, **G** & **H**. Results are averaged over 100 realisations, and all parameter values not specified here are listed in Table 1.

Discussion

Coevolution between effector cells and their cancer targets sets the stage for the emergence and subsequent development of a tumour. Based on expansive literature in this field [11, 31, 33], we focus on important perspectives which have not yet been addressed by the previous models: in particular, the stochastic nature of these early-stage small cell populations as well as explicit interactions between cancer cells and immune cells. Here, we model cancer-immune coevolution, wherein specialised effector cells respond to the presence of randomly-accumulated neoantigens in a growing cancer population. We uncover a variety of cancer-immune population dynamics, from the escape or extinction of the tumour to cycles characteristic of antagonistic interactions. We find that the suppression of the tumour by the immune system depends strongly on the cancer-immune interaction parameters, as well as rates of antigenic mutation accumulation in the cancer population. Using mutational distributions, we identify selection and the strength of cancer-immune interactions in the system, arguing for the importance of integrating population- and single cell-level data, especially in the context of informing immunotherapeutic practices with model predictions.

Instead of encapsulating the immune impact into a selection parameter, which assumes the effector population reacts immediately and perfectly to any new antigenic mutation [11, 36], we model the explicit interactions between cancer and immune cells. Our model unveils effector population dynamics during burst-like responses to growing subclones of the cancer population that possess antigenic mutations with high immunogenicity (see Figure 2). These immune population spikes serve to eliminate specific mutations from the cancer population. Via this selective killing of cancer cells, a process known as immunoselection [15], the immune system moulds the genetic landscape of the tumour in ways that are identifiable via sequencing data. For instance, the decrease in average antigenic mutations per cell (see Figure 4) and the truncation of the high-burden tail of the MBD (see Figure 5) demonstrate that cells with more mutations face stronger negative selection by immune response and are thus pruned from the population. It should be noted, however, that for certain cancers the neoantigenic landscape has been found to be only weakly impacted by cancer-immune coevolution [10]. Central to this modelling challenge is understanding the mutational process itself.

A fraction of somatic mutations arising in cancer populations give rise to the presentation of neoantigens [58]. This is a random process, wherein high mutational loads do not necessarily correspond to high antigenicities [15]. This implies that careful consideration of the mutational burdens of cells as well as the antigenicities of individual mutations is crucial to understanding the resulting evolutionary dynamics of the system. The total mutational burden of the tumour, however, is not a sufficient predictor of response to treatment unless mutations that have escaped are taken into consideration [12]. While the SFS and the single-cell MBD can inform and help quantify selection (see Figures 5 and A8), more work needs to be done to understand the impact of different mechanisms of immune escape on genetic data.

We demonstrate in Section A1 of the Appendix that a higher average SFS increases the growth of effector cells: indeed, the smaller the net impact on the effector population of cancer-immune interactions (that is, the average of $\alpha_i - \gamma_i$), the larger the mean of the SFS must be to support a growing effector population (see equation (A5)). Higher intratumour heterogeneity, captured by a higher mean of the SFS, not only translates into a more effective immune response, but also to better outcomes when treated with immunotherapy [36].

Informing treatment is one of the principal tenets of mathematical modelling in oncology [32, 59]. Some immunotherapies increase the ability of effector cells to kill cancer cells [15], while others, termed immune checkpoint inhibitors, reactivate immune predation in the case of antigenic mutations having escaped detection [12]. The neoantigens accumulated after escape thus work against the cell once immunotherapy renders the cell visible to the immune system once more, though immune evasion may still impede immunotherapy [33]. If, however, the antigenicities of the cancer cells are low due to immunoselection, the tumour will be less likely to respond well to immune checkpoint inhibition [12].

Few models have explored the relative advantages of different changes in tumour-immune interactions, which represent the impact of immunotherapies discussed above [45]. Wilkie & Hahnfeldt, for instance, demonstrated that resistance to immune predation plays a smaller role than effector recruitment [16]. Our results show that these relative advantages are highly dependent on the system itself: in Figure 3D, moving upwards from point c (decreasing the resistance to predation) has little impact on outcome, whereas moving to the right (increasing recruitment) changes the outcome drastically. Conversely, at point e, we notice the opposite effect: a change in predation resistance impacts the outcome but a change in recruitment does not. Importantly, only treatments that transform system parameters can succeed in a robust fashion, since only changing the state will still result in the same equilibria as before [59].

Limitations exist when trying to model the cancer-immune system. When the (antigenic) mutation rate is low, the fraction of the tumour visible to the immune system is too (see Figure A6). The cancer population dynamics are therefore largely neutral [11], though our model reveals complex effector dynamics. One can also assume a certain antigenic proportion in the tumour before immune recognition [36], or address the growth-threshold conjecture, which states that the immune system will respond to a large enough tumour growth rate, rather than a certain tumour size [33, 60]. The situation can also be further complicated by explicitly considering the composite state of an effector cell in the process of killing its cancer prey as a new conjugate type in the model, as has recently been done by Yang *et al.*, who demonstrated its possible impact on the resulting dynamics, including on the outcome and its time scale [61].

By employing an individual-based model, we can compare expected mutational distributions with corresponding genetic data. The signatures of selection and strength of cancer-immune interactions in the system along with a mechanistic knowledge of these interactions may help inform us of a tumour's evolutionary history, along with its immunotherapeutic potential.

Code Availability

Code of implementation of our model is available at https://github.com/Bio421/CancerImmuneCoevo under the MIT license.

Acknowledgements

L.W. was supported by the International Program for Candidates, Sun Yat-Sen University. C.M. was supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie EvoGamesPlus grant number 955708.

Competing interests

The authors declare no competing interests.

References

[1] Douglas Hanahan. "Hallmarks of cancer: New dimensions". In: *Cancer Discovery* 12.1 (2022), pp. 31–46.

[2] Antoine M. Dujon et al. "Identifying key questions in the ecology and evolution of cancer". In: *Evolutionary Applications* 14.4 (2021), pp. 877–892.

[3] Irina Kareva et al. "Predator-prey in tumor-immune interactions: A wrong model or just an incomplete one?" In: *Frontiers in Immunology* 12 (2021), p. 668221.

[4] Jason T George and Herbert Levine. "Sustained coevolution in a stochastic model of cancer-immune interaction". In: *Cancer Research* 80.4 (2020), pp. 811–819.

[5] Gavin P Dunn, Lloyd J Old, and Robert D Schreiber. "The three Es of cancer immunoediting". In: *Annual Review of Immunology* 22.1 (2004), pp. 329–360.

[6] Macfarlane Burnet. "Cancer—a biological approach". In: *British Medical Journal* 1.5022 (1957), p. 779.

[7] FM Burnet. "Immunological aspects of malignant disease". In: *The Lancet* 289.7501 (1967), pp. 1171–1174.

[8] Cristian Tomasetti and Bert Vogelstein. "Variation in cancer risk among tissues can be explained by the number of stem cell divisions". In: *Science* 347.6217 (2015), pp. 78–81.

[9] Catherine M Koebel et al. "Adaptive immunity maintains occult cancer in an equilibrium state". In: *Nature* 450.7171 (2007), pp. 903–907.

[10] Eszter Lakatos et al. "Epigenome and early selection determine the tumour-immune evolutionary trajectory of colorectal cancer". In: *bioRxiv* (2024).

[11] Eszter Lakatos et al. "Evolutionary dynamics of neoantigens in growing tumors". In: *Nature Genetics* 52.10 (2020), pp. 1057–1066.

[12] Luis Zapata et al. "Immune selection determines tumor antigenicity and influences response to checkpoint inhibitors". In: *Nature Genetics* 55.3 (2023), pp. 451–460.

[13] Ludmil B Alexandrov et al. "Signatures of mutational processes in human cancer". In: *Nature* 500.7463 (2013), pp. 415–421.

[14] Nicholas McGranahan et al. "Allele-specific HLA loss and immune escape in lung cancer evolution". In: *Cell* 171.6 (2017), pp. 1259–1271.

[15] Ton N Schumacher and Robert D Schreiber. "Neoantigens in cancer immunotherapy". In: *Science* 348.6230 (2015), pp. 69–74.

[16] Kathleen P Wilkie and Philip Hahnfeldt. "Mathematical models of immune-induced cancer dormancy and the emergence of immune evasion". In: *Interface Focus* 3.4 (2013), p. 20130010.

[17] Marta Łuksza et al. "A neoantigen fitness model predicts tumour response to checkpoint blockade immunotherapy". In: *Nature* 551.7681 (2017), pp. 517–520.

[18] Jimmy Van den Eynden et al. "Lack of detectable neoantigen depletion signals in the untreated cancer genome". In: *Nature Genetics* 51.12 (2019), pp. 1741–1748.

[19] Daniel S Chen and Ira Mellman. "Oncology meets immunology: The cancer-immunity cycle". In: *Immunity* 39.1 (2013), pp. 1–10.

[20] Dass S. Vinay et al. "Immune evasion in cancer: Mechanistic basis and therapeutic strategies". In: *Seminars in Cancer Biology*. Vol. 35. Elsevier. 2015, S185–S198.

[21] Emma Leschiera. "Mathematical models of tumour-immune interactions: Discrete and continuum approaches". PhD thesis. Sorbonne Université, UPMC University of Paris 6, 2022.

[22] Julieann Puleo and Kornelia Polyak. "A Darwinian perspective on tumor immune evasion". In: *Biochimica et Biophysica Acta (BBA) - Reviews on Cancer* 1877.1 (2022), p. 188671.

[23] Oscar Sotolongo-Costa et al. "Behavior of tumors under nonstationary therapy". In: *Physica D: Nonlinear Phenomena* 178.3-4 (2003), pp. 242–253.

[24] Luke Alexander Heirene et al. "Understanding Antibody-Target Antigen Interactions and the Avidity Effect Using Mathematical Modelling". In: *bioRxiv* (2024), pp. 2024–05.

[25] Charles J Mode. "A mathematical model for the co-evolution of obligate parasites and their hosts". In: *Evolution* (1958), pp. 158–165.

[26] Steve Paterson et al. "Antagonistic coevolution accelerates molecular evolution". In: *Nature* 464.7286 (2010), pp. 275–278.

[27] Masato Yamamichi and Stephen P. Ellner. "Antagonistic coevolution between quantitative and Mendelian traits". In: *Proceedings of the Royal Society B: Biological Sciences* 283.1827 (2016), p. 20152926.

[28] Weini Huang et al. "Dynamical trade-offs arise from antagonistic coevolution and decrease intraspecific diversity". In: *Nature Communications* 8.1 (2017), p. 2059.

[29] Raluca Eftimie, Jonathan L Bramson, and David JD Earn. "Interactions between the immune system and cancer: A brief review of non-spatial mathematical models". In: *Bulletin of Mathematical Biology* 73 (2011), pp. 2–32.

[30] Phineas T. Hamilton, Bradley R. Anholt, and Brad H. Nelson. "Tumour immunotherapy: lessons from predator–prey theory". In: *Nature Reviews Immunology* 22.12 (2022), pp. 765–775.

[31] Alberto d'Onofrio. "A general framework for modeling tumor-immune system competition and immunotherapy: Mathematical analysis and biomedical inferences". In: *Physica D: Nonlinear Phenomena* 208.3-4 (2005), pp. 220–235.

[32] Ivana Bozic, Alanna Sholokova, and Kamran Kaveh. "Neoantigen evolution and response to checkpoint inhibitor immunotherapy in colorectal cancer". In: *ResearchSquare preprint* (2024).

[33] Jason T George and Herbert Levine. "Stochastic modeling of tumor progression and immune evasion". In: *Journal of Theoretical Biology* 458 (2018), pp. 148–155.

[34] Jason T George and Herbert Levine. "Implications of tumor-immune coevolution on cancer evasion and optimized immunotherapy". In: *Trends in Cancer* 7.4 (2021), pp. 373–383.

[35] Jason T George and Herbert Levine. "Optimal cancer evasion in a dynamic immune microenvironment generates diverse post-escape tumor antigenicity profiles". In: *eLife* 12 (2023), e82786.

[36] Shaoqing Chen et al. "Frequency-dependent selection of neoantigens fosters tumor immune escape and predicts immunotherapy response". In: *Communications Biology* 7.1 (2024), p. 770.

[37] John N Thompson and Jeremy J Burdon. "Gene-for-gene coevolution between plants and parasites". In: *Nature* 360.6400 (1992), pp. 121–125.

[38] Andrei Korobeinikov, Konstantin E Starkov, and Paul A Valle. "Modeling cancer evolution: Evolutionary escape under immune system control". In: *Journal of Physics: Conference Series*. Vol. 811. 1. IOP Publishing. 2017, p. 012004.

[39] John A Adam and Nicola Bellomo. *A Survey of Models for Tumor-Immune System Dynamics*. Springer Science & Business Media, 2012.

[40] Daniel T Gillespie. "A general method for numerically simulating the stochastic time evolution of coupled chemical reactions". In: *Journal of Computational Physics* 22.4 (1976), pp. 403–434.

[41] Joshua Lederberg. "Replica plating and indirect selection of bacterial mutants: Isolation of preadaptive mutants in bacteria by sib selection". In: *Genetics* 121.3 (1989), p. 395.

[42] Motoo Kimura. "The number of heterozygous nucleotide sites maintained in a finite population due to steady flux of mutations". In: *Genetics* 61.4 (1969), pp. 893–903.

[43] Martina Baar et al. "A stochastic model for immunotherapy of cancer". In: *Scientific Reports* 6.1 (2016), p. 24169.

[44] Vladimir A Kuznetsov et al. "Nonlinear dynamics of immunogenic tumors: parameter estimation and global bifurcation analysis". In: *Bulletin of Mathematical Biology* 56.2 (1994), pp. 295–321.

[45] David Morselli et al. "A hybrid discrete-continuum modelling approach for the interactions of the immune system with oncolytic viral infections". In: *arXiv:2404.06459* (2024).

[46] Frank Nani and Herb I Freedman. "A mathematical model of cancer treatment by immunotherapy". In: *Mathematical Biosciences* 163.2 (2000), pp. 159–199.

[47] Crawford S Holling. "The components of predation as revealed by a study of small-mammal predation of the European Pine Sawfly". In: *The Canadian Entomologist* 91.5 (1959), pp. 293–320.

[48] Vladimir A Kuznetsov and Gary D Knott. "Modeling tumor regrowth and immunotherapy". In: *Mathematical and Computer Modelling* 33.12-13 (2001), pp. 1275–1287.

[49] Christo Morison, Dudley Stark, and Weini Huang. "Single-cell mutational burden distributions in birth-death processes". In: *arXiv preprint arXiv:2309.06355* (2023).

[50] Laura E Jones and Stephen P Ellner. "Effects of rapid prey evolution on predator–prey cycles". In: *Journal of Mathematical Biology* 55 (2007), pp. 541–573.

[51] Andrei Papkou et al. "Host–parasite coevolution: Why changing population size matters". In: *Zoology* 119.4 (2016), pp. 330–338.

[52] Rick Durrett. "Population genetics of neutral mutations in exponentially growing cancer cell populations". In: *The Annals of Applied Probability* 23.1 (2013), p. 230.

[53] Ivana Bozic, Jeffrey M Gerold, and Martin A Nowak. "Quantifying clonal and subclonal passenger mutations in cancer evolution". In: *PLOS Computational Biology* 12.2 (2016), e1004731.

[54] ML Samuels. "Distribution of the branching-process population among generations". In: *Journal of Applied Probability* 8.4 (1971), pp. 655–667.

[55] David Cheek and Samuel GG Johnston. "Ancestral reproductive bias in branching processes". In: *Journal of Mathematical Biology* 86.5 (2023), p. 70.

[56] Marius E Moeller et al. "Measures of genetic diversification in somatic tissues at bulk and single-cell resolution". In: *eLife* 12 (2024), RP89780.

[57] Einar Bjarki Gunnarsson, Kevin Leder, and Jasmine Foo. "Exact site frequency spectra of neutrally evolving tumors: A transition between power laws reveals a signature of cell viability". In: *Theoretical Population Biology* 142 (2021), pp. 67–90.

[58] Carsten Linnemann et al. "High-throughput epitope discovery reveals frequent recognition of neo-antigens by CD4+ T cells in human melanoma". In: *Nature Medicine* 21.1 (2015), pp. 81–85.

[59] Robert A Gatenby and Thomas L Vincent. "Application of quantitative models from population biology and evolutionary game theory to tumor therapeutic strategies". In: *Molecular Cancer Therapeutics* 2.9 (2003), pp. 919–927.

[60] Zvi Grossman and William E Paul. "Adaptive cellular interactions in the immune system: The tunable activation threshold and the significance of subthreshold responses." In: *Proceedings of the National Academy of Sciences* 89.21 (1992), pp. 10365–10369.

[61] Qianci Yang, Arne Traulsen, and Philipp Altrock. "Integration of immune cell-target cell conjugate dynamics changes the time scale of immune control of cancer". In: *bioRxiv* (2024).

[62] Anna Kane Laird. "Dynamics of tumour growth". In: *British Journal of Cancer* 18.3 (1964), p. 490.

[63] BP Kharlamov. "On the generation numbers of particles in a branching process with overlapping generations". In: *Theory of Probability & Its Applications* 14.1 (1969), pp. 44–50.

[64] Clemente F Arias et al. "The growth threshold conjecture: A theoretical framework for understanding T-cell tolerance". In: *Royal Society Open Science* 2.7 (2015), p. 150016.

Appendix

This appendix is split into three sections: deterministic analysis, stochastic analysis and extended figures. We first state the ODE limit of our stochastic model and draw some conclusions about intratumour heterogeneity and its impact on effector cell population growth (Section A1). We then inspect the first mutant clone to arise in the cancer cell population (Section A2), which suggests investigating the two regimes of effector cell population growth when passive and active recruitment dominate, respectively (Sections A3 and A4). We return to the stochastic model to state theoretical distributions from the literature (Section A5). Second, we define the neoantigen burden distribution and derive its expected distribution, as well as plot the correlation between the mean of the NBD and the mean of the SFS (Section A6). We then discuss early dynamics (Section A7). Next, we investigate the effector cell population spikes and derive approximations for their expected frequency and amplitudes (Section A8). Finally, we discuss the implications of a faster immune response due to an existing effector pool in the body on the genetic evidence for selection A9. Additional figures of stochastic simulations are found in Sections A10-A12.

Deterministic analysis

A1   ODE model

Associate to the antigenic mutation labelled i an effector population E_i with passive recruitment rate B and per-capita death rate D, along with $\alpha_i - \gamma_i$ the factor denoting the net difference between active recruitment and inhibition/exhaustion interactions between effector cells of type i and cancer cells possessing mutation i, whose population we denote C_i. The ODE describing the population dynamics of E_i is thus

$$\frac{\mathrm{d}E_i}{\mathrm{d}t} = B + E_i\left((\alpha_i - \gamma_i)\,C_i - D\right). \tag{A1}$$

We will write $E = \sum_i E_i$ for the total effector population. We will ignore neutral mutations in the ODE model, so unless otherwise specified, mutations are now assumed to be antigenic.

Because we are considering the emergence of cancer (that is, our populations are small), we model tumour growth exponentially. It is expected that the immune system is able to handle small tumours [5], so the tumour per-capita growth rate should, in principle, be bounded near populations of zero [31]. Note that the oft-used Gompertz model of tumour growth, which is unbounded at zero, was originally used for tumours whose initial sizes were on the order of thousands of cells [62].

Recall the following definitions from the main text. Let the elements of the site frequency spectrum (SFS) S_j denote the number of mutations occurring j times in the (cancer cell) population, so that $M = \sum_j S_j$ is the total number of mutations. Let the elements of the single-cell mutational burden distribution (MBD) B_k denote the number of cells possessing k mutations, so that $C = \sum_k B_k$ is the total cancer population. The total number of mutational occurrences (see the Supplementary Information of [49] for details) is then

$$U = \sum_{j=1}^{C} jS_j = \sum_{k=1}^{M} kB_k, \tag{A2}$$

so that the average appearances of a mutation (that is, the mean of the SFS) is U/M, and the average single-cell mutational burden (that is, the mean of the MBD) is U/N.

Notice that for a fixed mutation i, when $\alpha_i - \gamma_i$ is high, we expect C_i to be low, since more effector cells of type i are around (having been recruited but not inhibited nor exhausted) to deplete the population C_i. Thus we expect that $\mathrm{cov}(\alpha_i - \gamma_i, C_i) \leq 0$ for all i.

Hoping to inspect the overall effector population, we sum (A1) over mutations $1 \leq i \leq M$ to find

$$\frac{\mathrm{d}E}{\mathrm{d}t} = MB - DE + \langle(\alpha_i - \gamma_i)C_i\rangle E, \tag{A3}$$

for $\langle\cdot\rangle$ the average over i. Now,

$$\langle(\alpha_i - \gamma_i)C_i\rangle = \langle\alpha_i - \gamma_i\rangle\langle C_i\rangle + \mathrm{cov}(\alpha_i - \gamma_i, C_i) \leq \langle\alpha_i - \gamma_i\rangle\langle C_i\rangle = \Delta\frac{U}{M},$$

where we have written $\Delta = \langle \alpha_i - \gamma_i \rangle$ and identified that $\langle C_i \rangle$ is the average number of cells with a given mutation, which is the same as the average appearances of a mutation, or U/M. Thus, (A3) becomes

$$\frac{\mathrm{d}E}{\mathrm{d}t} \leq MB + \left(\frac{\Delta U}{M} - D \right) E. \tag{A4}$$

Ignore passive recruitment (that is, the MB term in (A4)) for the moment; we will later call this the active recruitment regime, where interactions with target cancer cells contribute more to the effector population than the drift term B. A conservative condition for the effector population to be growing, $\mathrm{d}E/\mathrm{d}t > 0$, is here

$$\frac{U}{M} > \frac{D}{\Delta}. \tag{A5}$$

This inequality can be interpreted in the following way: if the net impact of effector-cancer interactions Δ (assumed to be positive, as otherwise the effector population always decreases, barring passive recruitment effects which we assume to be small) is small, the SFS must be heavy-tailed (that is, more heterogeneous) for the effector population size to grow.

A2   The first mutant clone

Suppose an antigenic mutation arises with antigenicity A_1 and immunogenicity I_1 in a single progenitor cell within a population of neutral cancer cells. We want to inspect the ensuing clone with population C_1 and the corresponding effector population E_1. Write b and d for the per-capita birth and (intrinsic) death rates (where we incorporate the immune escape rate into d for simplicity) of the cancer cells, regardless of mutational load, and the net growth rate (which we will always consider to be positive) $r = b - d > 0$. Now, the clone's population satisfies

$$\frac{\mathrm{d}C_1}{\mathrm{d}t} = rC_1 - \beta_1 E_1 C_1, \tag{A6}$$

where we have written $\beta_1 = \beta_0 A_1$ for the killing rate of the cancer-effector interactions. From (A1), the effector population E_1 satisfies

$$\frac{\mathrm{d}E_1}{\mathrm{d}t} = B_1 + (\alpha_1 - \gamma_1)C_1 E_1 - D_1 E_1. \tag{A7}$$

Label the effector population where the dominant recruitment mode transitions from passive to active as $E_1 = E^*$. If we consider the coarse approximation of E^* being a strict transition (that is, where all recruitment when $E_1 < E^*$ is passive and all recruitment when $E_1 > E^*$ is active; note that we will almost always be looking at populations far from E^*, so we will not need to consider populations $E_1 \approx E^*$), then we are interested in two time regimes: $t \ll E^*/B_1$ (the time at which we expect E^* cells to have been recruited passively) and $t \gg E^*/B_1$.

If the mutation arises at time $t = 0$, it takes $1/B_1$ for the first effector cell to appear. At this point, in the deterministic framework, there are $C_1(t = 1/B_1) = \mathrm{e}^{r/B_1}$ cancer cells.

A3   Passive recruitment regime

For $t \ll E^*/B_1$ we have $E_1 = B_1 t$, which allows us to solve (A6) for $C_1(t)$:

$$\frac{\mathrm{d}C_1}{\mathrm{d}t} = (r - \beta_1 B_1 t) C_1 \quad \Rightarrow \quad C_1(t) = \exp\left(rt - \frac{1}{2}\beta_1 B_1 t^2 \right) \leq \mathrm{e}^{rt}. \tag{A8}$$

We also know that the cancer cell population at the transition time will be bounded from above:

$$C_1\left(t \approx \frac{E^*}{B_1} \right) \leq \exp\left(\frac{rE^*}{B_1} \right), \tag{A9}$$

since the right-hand side of (A9) is the population supposing the effector cells have not killed any of the cancer cells.

Recall from (A7) that the transition between recruitment types occurs when $B_1 = \Delta_1 E_1 C_1$ (where $\Delta_1 = \alpha_1 - \gamma_1$), which gives an equation for E^* that can be numerically (though not analytically) solved:

$$B_1 = \Delta_1 E^* \mathrm{e}^{rE^*/B_1}. \tag{A10}$$

A4 Active recruitment regime

When $t \gg E^*/B_1$, we ignore the B_1 term of (A7) and are thus left with what appear to be predator-prey equations for E_1 and C_1, as long as Δ_1 is constant in C_1 (which we will see is not necessarily the case). Because the active recruitment α_1 is a type-II functional response rather than a type-I bilinear response (like the inhibition/exhaustion term γ_1), as shown in Figure A1, we should subdivide this regime further into the intervals delineated by the crossover points $C_1 = C^*$ and $C_1 = C^{**}$.

Figure A1. The net difference Δ_1 (solid pink line) between active recruitment α_1 (dotted blue line) and inhibition/exhaustion γ_1 (dashed red line) interactions between effector and cancer cells splits the cancer cell population C_1 into three natural regions: $[0, C^*)$, $[C^*, C^{**})$ and $[C^{**}, \infty)$. Note that the active recruitment α_1 is a type-II functional response (reaching a maximum of h_α^{-1} as $C_1 \to \infty$), whose shape more resembles the dotted pale blue line than our piecewise linear approximation shown with a thicker line; this also translates to the thinner solid pink line for $\Delta_1 = \alpha_1 - \gamma_1$.

When $C_1 > C^{**}$, $\Delta_1 < 0$, so the effector population will decrease no matter what, barring passive recruitment. This means that if a subclone of cancer cells with a certain antigenic mutation can grow to a specific density, then it will successfully evade this immune response, since its corresponding effector population will not be able to grow to keep up with it. This, however, is restricted only to the effector type corresponding to the mutation that originated the clone; other antigenic mutations will arise in subclones, and therefore the clone overall will not have completely evaded the immune system.

When $C_1 < C^*$ (especially for $C_1 \ll C^*$, when α_1 is well approximated by a linear function), we have a constant of motion from the predator-prey system:

$$V = \Delta_1 C_1 - D_1 \log C_1 + \beta_1 E_1 - r \log E_1. \tag{A11}$$

(Though since the passive recruitment contributes a drift term, V is not identically constant.) From standard linear stability analysis we also have equilibrium points at $(C_1, E_1) = (0, 0)$ and $(D_1/\Delta_1, r/\beta_1)$. Beyond the drift term that increases the population of the effectors (notably at the latter equilibrium point of the simplified system), the aforementioned effect of new antigenic mutations within the C_1 clone giving rise to other predating effector populations is also obscured by our simplifications. These both contribute to decrease the population C_1.

The opposite effect is seem in the interval $C^* < C_1 < C^{**}$, when Δ_1 is positive but decreasing: the effector population E_1 may still grow, but a positive contribution to the equilibrium population of cancer cells exists in the form of the decreasing effectiveness of the effector cells.

Stochastic analysis

A5 Expected mutational distributions

We begin by stating known expressions for the expected SFS and MBD in the case of neutral evolution, which we used for plots in Figures 4 and 5. From [57], the expected neutral SFS at large population size K with mutations accumulating at rate μ, conditional on survival (which we do by filtering out the realisations that do not make it to population K) is given by

$$\mathbb{E}\left[S_j \,|\, \text{survival}\right] \approx \mu K \sum_{i=0}^{\infty} \frac{(d/b)^i}{(i+j)(i+j+1)}. \tag{A12}$$

We take $\mu = \lambda(1 - p_a)$ for neutral mutations, and $\mu = \lambda p_a$ for antigenic mutations.

Under the same assumptions, the expected neutral MBD can be calculated as follows [49]:

$$\mathbb{E}\left[B_k \,|\, \text{survival}\right] \approx \sum_{\ell} D_\ell \frac{\mathrm{e}^{-\ell\mu}(\ell\mu)^k}{k!} \approx \sum_{\ell} \left(\frac{(b+d)\mu\log K}{b-d}\right)^\ell \frac{K^{-\mu(b+d)/(b-d)}}{\ell!} \frac{\mathrm{e}^{-\ell\mu}(\ell\mu)^k}{k!}, \tag{A13}$$

where the sum over ℓ spans at most the number of events (births or deaths), though the sum converges faster. We have used a continuous-time approximate version from [63] for the expected division distribution D_ℓ instead of the discrete-time expression from [49] for ease of computation, along with the approximation that the population grows to K in $\log K/(b-d)$ time, which arises by inspection of the deterministic case of exponential growth.

A6 Neoantigen burden distribution

The immune system reacts to the antigenicities of the antigenic mutations possessed by a cancer cell rather than simply its (antigenic) mutational burden. To capture this nuance, we define the neoantigen burden N_ℓ as the number of cells with cumulative antigenicity $\sum_i A_i$ (where the sum spans the antigenic mutations i within a cell) in $[\ell, \ell+1)$. Taken together, the elements N_ℓ form the neoantigen burden distribution (NBD).

Following the derivation of the MBD from the division distribution in [49], we write $A_{i,j,k} \sim \text{Exp}(1)$ for the antigenicity of the ith antigenic mutation of the jth cell (for some labelling of cells $1 \leq j \leq B_k$) having a mutational burden of k and $\mathbb{1}_A$ for the indicator function that is 1 on the set A and 0 otherwise. Conditioning on knowledge of the antigenic MBD $\{B_k\}_{k=1}^{M}$, we find

$$\mathbb{E}\left[N_\ell | \{B_k\}_k\right] = \mathbb{E}\left[\sum_{k=1}^{M}\sum_{j=1}^{B_k} \mathbb{1}_{\left\{\sum_{i=1}^{k} A_{i,j,k} \in [\ell, \ell+1)\right\}} \Bigg| \{B_k\}_k\right] = \sum_{k=1}^{M} B_k \mathbb{E}\left[\mathbb{1}_{\left\{\sum_{i=1}^{k} A_{i,j,k} \in [\ell, \ell+1)\right\}}\right]. \tag{A14}$$

A sum of k exponentially-distributed random variables with means ν obeys an Erlang distribution, which has probability density function

$$f(t; k, \nu) = \frac{\nu^k t^{k-1} \mathrm{e}^{-\nu t}}{(k-1)!}. \tag{A15}$$

We take expectations of both sides of equation (A14) (where $\nu = 1$) and use the law of total expectation $\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X|Y]]$ and equation (A15) to find

$$\mathbb{E}\left[N_\ell\right] = \sum_{k=1}^{M} \mathbb{E}\left[B_k\right] \int_{\ell}^{\ell+1} \frac{t^{k-1} \mathrm{e}^{-t}}{(k-1)!} \, \mathrm{d}t = \mathrm{e}^{-\ell} \sum_{k=1}^{M} \mathbb{E}\left[B_k\right] \sum_{k'=0}^{k-1} \frac{\ell^{k'} - \mathrm{e}^{-1}(\ell+1)^{k'}}{k'!}. \tag{A16}$$

This allows for a flexible definition of the NBD: we have chosen to discretise at integer intervals $[\ell, \ell+1)$, but could have kept it continuous with probability density function $\sum_k \mathbb{E}[B_k] t^{k-1}\mathrm{e}^{-t}/(k-1)!$.

Figure A2 depicts the average NBD (along with the corresponding construction for the immunogenicity) over 100 realisations, when $\lambda = 1$ (Figures A2A-B) and when $\lambda = 10$ (Figures A2C-D). Both low and high immune impact parameter sets are plotted, to demonstrate that much like the MBD (see Figures 5 and A8), the tail of the distribution is shortened in the case of higher selection. This is because cells with higher antigenicities (likewise for those with higher immunogenicities) face stronger negative selection from the effector population.

Figure A2. Neoantigen burden distribution (NBD) (yellow lines) and the corresponding distribution for immunogenicity (purple lines) averaged over 100 realisations, when $\lambda = 1$ (**A**-**B**) and $\lambda = 10$ (**C**-**D**), for both low (dotted lines) and high (dashed lines) immune effectiveness. All parameter values not specified here are listed in Table 1 of the main text.

A7 Early stochastic dynamics

Similar to the deterministic discussion of Sections A3 and A4, passive and active recruitment regimes can also be identified in the stochastic model of the main text. On average, an effector population of type $i = 1$ will have passive capacity B/D: that is, at a population of $E_1 = B/D$, its passive recruitment rate and death rate are equal, and thus ignoring its interactions with its target population C_1, E_1 will hover around this value. Once C_1 grows to $D/(\alpha_1 - \gamma_1)$, however, the active recruitment rate will equal the death rate. We can thus qualitatively describe the early dynamics of a cancer clone as follows: the mutation arises; the effector population grows to B/D; the clone grows to $D/(\alpha_1 - \gamma_1)$, at which point the effector population increases to properly (that is, via active recruitment) combat the threat.

For most parameter choices in the main text, $B/D = O(1)$ and $D/(\alpha_1 - \gamma_1) = O(10^2)$, so that the effector population only manages to grow to more than a few cells once its targets number around a hundred. (It is also worth mentioning that early ODE models of tumour-immune dynamics phrased in a predator-prey fashion use Kuznetsov *et al.*'s estimate of effector passive recruitment being $O(10^4)$ cells per day [44]. With $O(0.1)$ or $O(1)$ new antigenic mutations (and thus effector types) generated in each division, when our cancer cell population is $O(10^4)$, we therefore expect a passive recruitment similar to Kuznetsov *et al.* [44].)

Models of immune responses vary in their triggers: Chen & Xie *et al.* suppose a minimal proportion of antigenic cancer cells [36], the growth threshold conjecture requires a minimal cancer growth rate [64], and most ODE models as well as Lakatos *et al.* suppose instantaneous responses [11, 31]. While our mechanistic model does not assume frequency dependence, the previous discussion implies that the immune response implicitly requires a certain cancer density before it can grow effectively to fight.

A8   Effector cell population spikes

As identified in individual-based simulations (see Figure 2, for instance), effector subpopulations in the stochastic model can undergo rapid spiking events, before vanishing to zero when the targeted mutant (labelled by i) goes extinct (and thus the effector cells have no more prey and so are removed from the system). These spikes arise when the mutation has immunogenicity I_i much higher than the mean $\langle I \rangle$, as well as an antigenicity A_i lower than the mean $\langle A \rangle$, as otherwise the spike would not need to grow since the effector cells would be more efficient at killing. Qualitatively, we will write $\hat{I}$ ($\hat{A}$) for the threshold immunogenicity (antigenicity) and above (below) which spikes may occur.

Returning to the language of predator-prey dynamics discussed in Section A4, effector spikes arise when the populations enter a cycle that ends when $C_i(t) < 1$ (that is, extinction of mutation i). We write $i = 1$ for the mutation causes the spike, since we consider it to be arising in a neutral cell. (We argue qualitatively that were other antigenic mutations carried by the cell to generate an effector spike, they already would have; thus we expect most spikes to arise in neutral or almost neutral cells.) Now, recalling the constant of motion (A11) and setting $C_1 = 1$, we obtain an equation for E_1 (as a function of $\alpha_1 - \gamma_1$), the approximate (due to the stochasticity of the model as well as the simplifying assumptions above) amplitude of the spike:

$$V = \alpha_1 - \gamma_1 + \beta_1 E_1 - r \log E_1, \tag{A17}$$

which can have up to two solutions and can be numerically solved. Differentiating (A17) implicitly with respect to I_1, we obtain an ODE for E_1:

$$\frac{\mathrm{d}E_1}{\mathrm{d}I_1} = \frac{\mathrm{d}E_1}{\mathrm{d}(\alpha_1 - \gamma_1)} \frac{\mathrm{d}(\alpha_1 - \gamma_1)}{\mathrm{d}I_1} = \frac{E_1}{r - \beta_1 E_1} \frac{\mathrm{d}\alpha_1}{\mathrm{d}I_1}. \tag{A18}$$

When the cancer clone is decreasing in population, the first term is negative; the second term is always negative, since α_1 is increasing in I_1 by inspection of (1). Thus, we have that $\mathrm{d}E_1/\mathrm{d}I_1 > 0$, confirming that a higher immunogenicity does increase the corresponding effector population, as expected.

Using that immunogenicities and antigenicities are drawn from exponential distributions with means 1, the rate of spikes r_{spike} is then

$$r_{\mathrm{spike}} = p_a \lambda \mathrm{e}^{-(\hat{I}+\hat{A})}, \tag{A19}$$

since the rate of new antigenic mutations and the probabilities $\mathbb{P}(I_i > \hat{I})$ and $\mathbb{P}(A_i > \hat{A})$ are independent.

A9   Faster immune response

Our immune system possesses an innate pool of effector cells, able to respond to threats without the time-consuming training assumed implicitly in our specialised model of active and passive recruitment [19]. With some probability, this pool might include effector cells that can recognise newly arising neoantigens quickly. Thus, rather than having to wait for the first passively-recruited effector cell, the body could react faster to novel antigenic mutations.

To investigate the impact of this mechanism on the observed coevolution, we modified our model to include an effector cell of type i upon initiation of antigenic mutation i. Figures A3A-B show the average number of antigenic mutations per cell and Figures A3C-D show the SFS, in the case of $\lambda = 1$.

When comparing with the corresponding figures for the model used throughout the main text (Figures 4A-B, 5B and 5F, respectively), the patterns appear the same, though more evident. This reinforces that the speed of the immune response increases its effectiveness though does not change the qualitative nature of cancer-immune coevolution. However, rather than having a guaranteed effector cell that can handle a newly-arising neoantigen, there is some probability of the effector pool containing the desired specialty. Thus we expect that in the body, selection will follow the patterns shown in both cases (Figures 4 and 5 versus Figure A3), at some magnitude between the two, related to the aforementioned probability.

Figure A3. Genetic evidence of selection for a faster-acting immune system. **A-B** Average number of antigenic mutations per cell in solid oranges lines for several representative realisations when $\lambda = 1$. Theoretical prediction for the accumulation of neutral mutations per cell in an exponentially-growing population shown in grey dashed line. **C-D** Simulated SFS for antigenic mutations averaged over 100 realisations when $\lambda = 10$ (orange points), along with the theoretical predictions (black dashed lines) in the absence of an immune response. **A** & **C** Low immune effect: $\alpha_0 = 0.002$ and $\beta_0 = 0.001$. **B** & **C** High immune effect: $\alpha_0 = 0.03$ and $\beta_0 = 0.3$. (Parameter sets chosen as points a and b from Figure 3A.)

Extended figures

A10 Population dynamics

Figure A4 shows representative population dynamics (first row), along with the corresponding composition of the tumour population (second row), when $\lambda = 1$ (Figures A4A-D) and $\lambda = 10$ (Figures A4E-J). In the upper row, a blue histogram depicts the end times for the realisations, along with the average time T_K to reach population size $K = 3 \times 10^4$, when applicable. (As seen in Figure A4I, all realisations go extinct, so no such T_K is defined.) The red lines in the second row refer to the number of allowed antigenic mutations possessed by a cancer cell to consider it neutral. That is, the line labelled by 3 means that cells with 3 antigenic mutations are deemed neutral, whereas those with 4 antigenic mutations are considered antigenic and thus contribute to the proportion of the tumour that is antigenic. Note that since the immunogenicities and antigenicities are drawn from exponential distributions with mode 0, it is likely that many purportedly antigenic mutations i have very weak immunogenicity $I_i \approx 0$ and antigenicity $A_i \approx 0$. In this case, they would interact very rarely with the immune system (being effectively neutral), while still contributing to a count of cancer cells carrying antigenic mutations.

In Figure A4, the stopping time for the realisations T_{end} was made large so to establish a good estimate on the average time T_K that the cancer cell population (in the non-suppressed case) reaches the stopping population size K. Due to the stochasticity of the model, there is a distribution over end times for each realisation even for a single set of parameters. It is clear from the difference in values of the mean T_K for Figures A4A and A4C (likewise for Figures A4E and A4G) that the immune system plays a role in slowing the cancer growth, since avg $T_K^{\mathbf{A}} <$ avg $T_K^{\mathbf{C}}$ (and avg $T_K^{\mathbf{E}} <$ avg $T_K^{\mathbf{G}}$).

A11 Heat maps

Figures A5A and A5B depict the extinction times for simulations run with $\lambda = 1$ and $\lambda = 10$, respectively. Figures A6A-C and A6D-F depicts what proportion of tumour cells carry antigenic mutations for simulations run with $\lambda = 1$ and $\lambda = 10$, respectively. The first row allows no antigenic mutations in neutral cancer cells; the second (third) row allows one (two) antigenic mutation(s) before a cell is considered antigenic.

A12 Genetics

When hyper-mutated tumours are considered ($\lambda = 10$), the genetic footprints of selection are more pronounced. Figure A7 depicts the average number of antigenic mutations per cancer cell and Figure A8 shows the SFS and MBD in this regime.

Figure A4. Cancer-effector population dynamics (first row) and tumour antigenicity (second row) for five different sets of parameters. **A-B** $\lambda = 1$, $\alpha_0 = 0.002$ and $\beta_0 = 0.001$. **C-D** $\lambda = 1$, $\alpha_0 = 0.03$ and $\beta_0 = 0.3$. **E-F** $\lambda = 10$, $\alpha_0 = 0.002$ and $\beta_0 = 0.001$. **G-H** $\lambda = 10$, $\alpha_0 = 0.005$ and $\beta_0 = 0.01$. **I-J** $\lambda = 10$, $\alpha_0 = 0.03$ and $\beta_0 = 0.3$. In the first row, red lines depict cancer cell populations for several representative realisations. Blue histograms above show the final time T_K where the population size reached $K = 3 \times 10^4$ for all non-extinct realisations, with a mean given by a dashed vertical black line. In the second row, red lines with increasing paleness, labelled by $i = 0$, 1, 3 or 5, depict the proportion of cancer cells contain i or more antigenic mutations. All parameter values not specified here are listed in Table 1 of the main text.

Figure A5. Extinction time heat maps for $\lambda = 1$ (**A**) and $\lambda = 10$ (**B**). All parameter values not specified here are listed in Table 1 of the main text.

Figure A6. Heat maps depicting the tumour composition (that is, the proportion of tumour cells that are antigenic) for $\lambda = 1$ (**A-C**) and $\lambda = 10$ (**D-F**). The first row is where one antigenic mutation held by a cell makes the cell antigenic; the second and third rows allow for one and two antigenic mutations (respectively) to be possessed by a cancer cell while still considering it neutral. Points a and c (respectively b and d) label parameter sets of low (respectively high) immune effectiveness. All parameter values not specified here are listed in Table 1 of the main text.

Figure A7. Average number of antigenic mutations in solid oranges lines for several representative realisations when $\lambda = 10$. Theoretical prediction for the accumulation of neutral mutations in an exponentially-growing population shown in grey dashed line. **A** Low immune effect: $\alpha_0 = 0.002$ and $\beta_0 = 0.001$. **B** Middling immune effect: $\alpha_0 = 0.005$ and $\beta_0 = 0.01$. **C** High immune effect: $\alpha_0 = 0.03$ and $\beta_0 = 0.3$. (Parameter sets for **A** and **C** chosen as points c and d from Figure 3C, respectively.)

Figure A8. Genetic markers of selection: SFS (**A-B** & **E-F**) and MBD (**C-D** & **G-H**) for low (left column; $\alpha_0 = 0.002$ and $\beta_0 = 0.001$) and high (right column; $\alpha_0 = 0.005$ and $\beta_0 = 0.01$) immune effectiveness averaged over 100 realisations when $\lambda = 10$, along with the theoretical predictions (black dashed lines) in the absence of an immune response. Green data represents neutral mutations (**A**, **C**, **E** & **G**) and orange data represents antigenic mutations (**B**, **D**, **F** & **H**), with dashed vertical lines representing the means of the distributions for MBDs in panels **C**, **D**, **G** & **H**. All parameter values not specified here are listed in Table 1 of the main text.