ODD Documentation for the General Model of FGM as a Social Norm of Coordination

1. Purpose

The 'general model of FGM as a social norm of coordination' is an agent-based model of the social processes through which the practice of FGM persists, increases or decreases in localised populations. Specifically, it addresses social processes related to social coordination, where actors face incentives to match the behaviours of others. It represents 'everyday' processes of coordination in the absence of outside intervention, as well as the specific potential effects of targeted interventions designed to discourage FGM. The model, which can be seen as an expansive extension of formal models already proposed in the FGM policy literature, is designed to address a number of problems. Chapter 5 of the associated thesis identified a number of areas of potential uncertainty in the design of formal models of FGM as a norm of coordination. However, these explorations of individual features did not allow an assessment of the full range of dynamics that can occur when different areas of potential uncertainty are combined together. Neither did it allow an assessment of which of these areas of uncertainty area was most important (i.e. should be prioritised in calibration) when taking into account possible combinations of these features. The general model solves these problems by incorporating and parametrizing these different areas of uncertainty in a single model, for which previous models are special cases (i.e. particular parametrizations). By representing (and parametrizing) a global space of possible designs for coordination models of FGM, the corresponding space of possible predicted dynamics can be explored. This, in turn, makes it possible to characterise the relative importance of different areas of uncertainty in the model design, using techniques from global sensitivity analysis. Furthermore, the model can be applied directly within a *possibilistic failure scenario* context, to identify potential policy-failure scenarios – although these scenarios may have limited credibility in the absence of further empirical calibration. Empirical calibration of parts of the model has been carried out, and this is documented in the associated thesis (Chapter 7).

2. Entities, state variables and scales

The entities in the model are agents which represent individual social actors in FGM-practicing localised populations (typically territorial communities).

These entities are characterised by the following state-variables: the intrinsic value they place on FGM (which affects their willingness to participate or not participate in FGM-related activities - H_i), their authority (which affects the degree of pressure they exert on other actors to coordinate with them w_i), their autonomy (which affects the relation between social pressures experienced by the agent, and associated social costs, from their perspective: α_i), their status as a final decision-maker about girls in their household m_i, and their role as either tolerant of miscoordination by others ($\beta_{1i} = 0, \beta_{2i} = 0$), or an enforcer of FGM practice ($\beta_{1i} = 1$) or an enforcer of FGM abandonment (β_{2i}), both of which affect their contribution to norm-enforcement as a source of social influence.

Actors are also characterised by an *x* (x_i) and *y* (y_i) coordinate position in network-space (see below). Actors are further characterised by their *social, household* and *decision-maker* reference groups ($\{\text{social-ref}\}_i, \{\text{household}\}_i, \{\text{decision-maker}\}_\text{i}$), which are sets of other actors to whom they refer when assessing coordination incentives.

Actors' behavioural states are characterised by two dummy variables, one representing participation in FGM-related activities last time this arose (d_i, e.g. consenting to or helping plan/prepare cutting events, attending the event and/or aftermath, making a deal with circumcizors, paying circumcision fees, supporting the educational component of the ceremony, or declining to do these things, explicitly expressing intention not to support cutting, etc.) and another representing their final decision about whether a girl in their household should be cut (θ_i, the last time this arose, and assuming that they are a final decision-maker).

Actors are situated within a 'network space' which is a 2D coordinate space (the size is arbitrary – see section 8.3), that wraps on its horizontal axis. This space does not have a direct spatial analogue, but rather serves the instrumental function of constructing reference networks with a number of desirable properties- such as clustering, localization of social interaction, and homophily- which are controlled by global parameters of the environment.

The population of agents sits within a global environment, which is characterised by a large number of global parameters which represent different possibilities for the representation of FGM-related social coordination dynamics. In other words, the parameters in the global environment allow for a range of possible 'models' of dynamics and allow the modeller to explore a wide space of such possibilities. The space of global parameters can be seen as encompassing a certain amount of stochastic uncertainty (situations which may vary across different real-world communities) and epistemic uncertainty (possibilities whose adequacy as a model of the real situation is uncertain).

Global parameters control the basic initialization of the model by controlling the number of actors in the population (between 200 and 2000, n-actors), and the proportion of those actors who begin the simulation practicing FGM. They further control the *distribution* of heterogeneous characteristics across the agent population through parameterizing characteristics of the distributions of: the intrinsic value of FGM, autonomy, authority. This is achieved by parametrizing the mean and variance of underlying beta-distributions in each case (μ_H, μ_{alpha}, μ_w, σ_H^2, σ_α^2, σ_w^2).

The pairwise correlations between the intrinsic value of FGM, and actor's autonomy, authority and y-coordinate (each with intrinsic-value) are also parameterised (ρ_{Hw}, $\rho_{H\alpha}$, ρ_{Hy}). Further parameters control the proportion of actors selected as enforcers of FGM practice (s_6) or FGM abandonment (s_7) (conditional on their participation in FGM, or not) and the rules used to select which actors take on this role in each case (a_1, a_2, i.e. those with more authority, those with stronger beliefs about the intrinsic value or cost of FGM, etc.). Another parameter controls whether at initialization, participation in FGM is assigned at random, or whether actors who value FGM the most are selected first as initial participants (a_4).

Other global parameters control the decision-functions of actors by affecting the maximum possible intrinsic *value* attributed to FGM across the population of agents (relative to the maximum possible intrinsic *cost* - s_2), as well as: the maximum social miscoordination costs for abandoning the practice (relative to the maximum miscoordination costs for abandonment - s_1), the (potentially non-linear) rate with which social costs accumulate with the degree of social pressure (for FGM practise and non-practise: V, δ_1, δ_2), and the relative importance of influence from actors' three reference groups: *social*, *household* and *decision-maker* (s_3, s_4). A further parameter (s_5) controls whether social influence from the social and household reference groups is primarily implicit (depending on the weighted proportion of actors engaged in FGM participation) or primarily explicit (depending on the weighted proportion of actors engaged in pro-FGM *enforcement* or anti-FGM *enforcement*).

Other global parameters control the interaction structure of the population, by controlling: the average size of the actors' social reference groups (i.e. average degree in the social reference network, $\mu_{Rsocial}$), the average size of households μ_{hsize} (see below), the relative increase in the 'network reach' of the decision-maker reference groups relative to social reference groups (s_8), and the localization of household reference groups within the network space (which controls whether household membership tends to be homophilic and whether household networks span disparate parts of the social network by dividing the network space into η_1 discrete partitions and randomly creating household network cliques within these partitions).

Finally, the process flow of the model is controlled in terms of the number of coordination opportunities that occur per year (see time-scales below) and the order in which actors make decisions in each time step (e.g. random-sequential, simultaneous, high-authority actors first, etc.).

Simulated interventions are controlled by environmental variables which determine: the proportion of the population who are targeted (z_1), the strength of the educational component of the intervention (z_2), the minimum probability that actors who oppose FGM following intervention education will 'respond' positively (z_3) and be willing to form a coalition of non-practitioners with other actors, the number of other non-participating actors that responding-actors will try to recruit to the coalition (z_4), the proportion of those actors who were successfully recruited to the intervention that will join one-another's social reference group (z_5), the rules used to target actors for the intervention (z_6 - twelve options, including at random, high authority actors, actors within a localised part of the network, etc.) and the probability that intervention participants will become anti-FGM enforcers (z_7).

The temporal scale in the model is understood relative to the scale of a single year. A proportion of actors in the population make a decision about the cutting status of a girl in their household in a single year. This can be estimated from: the expected number of girls per-household at a given moment, from the number of households in the community and from the expectation that a decision must be made about all girls before they reach the age of 15 (see section 8.5). Against this fixed (relative to the population and households) rate of FGM cutting decisions, the number of decisions actors must take about FGM participation per year is controlled by a parameter (η_2), and can be between once and twelve times per year.

3. Process Overview and Scheduling

The initialization process involves the following sub-processes and is the first step in all uses of the model:

Table 1: Initialization Sub-Processes

SUB-PROCESS	BRIEF DESCRIPTION
Create base population	Create a population of N actors
Distribute population characteristics	Distribute actors' characteristics, including: their perceived intrinsic value of FGM (H_i), their authority (w_i), their autonomy (α_i) and their x and y-coordinates (y_i, x_i) according to the desired marginal ($\mu_H, \sigma_H^2, \mu_w, \sigma_w^2, \mu_\alpha, \sigma_\alpha^2$) and joints distributions of these characteristics ($\rho_{Hw}, \rho_{H\alpha}, \rho_{Hy}$). x_i and y_i always have a uniform distribution.
Set Spatial Location	Actors are placed in a spatial location in the 2D 'network space' according to their x_i and y_i coordination attributes.
Form Reference Networks	Actors construct social ($\{\text{social-ref}\}_i$), household ($\{\text{household}\}_i$), and decision-maker ($\{\text{decision-maker}\}_i$) reference groups. A random 1 + Poisson($\mu = 1$) actors in each household are assigned the role of 'decision-maker' ($m_i \rightarrow 1$).
Setup Initial Behaviors	Actors initial behaviours (practicing FGM or not), as well as their status as enforcers of the practice of FGM (β_1, β_2), are assigned according to the relevant assignment rules (a_1, a_2), including whether actors are initialised as practicing at random, or in descending order of FGM preference (a_3).

Following the initialization of the model, where applicable, a simulated intervention is immediately run. In this scenario, all actors are initialised as practicing FGM (and cutting girls in their household, etc.), removing some sensitivity of the model to initialization conditions.

The simulated intervention involves the following sub-processes:

Table 2: Simulated Intervention Sub-processes

SUB-PROCESS	**BRIEF DESCRIPTION**
Target actors for intervention (including educational effect)	A proportion of actors (z_1) are targeted for intervention participation, according to a preferential targeting rule (z_6 e.g. ascending order of preference). *'Educational Effect'* These actors receive the educational component of the intervention; their new (decreased) belief about the intrinsic value of FGM is a function of their previous belief (degree of change controlled by z_2). Agent's perception of the value of FGM *before* the educational effect of the intervention is defined: $$H_i = -\widehat{M} + (q \cdot [\widehat{M} + (\widehat{M} \cdot s_2)])$$ *After* the educational effect of the intervention, it is defined: $$H_i = -\widehat{M} + ([(1 - z_2) \cdot q] \cdot [\widehat{M} + (\widehat{M} \cdot s_2)])$$
Implement Intervention (Including Organised Diffusion)	With a probability that decreases with increased belief in the value of FGM (this probability is always zero if FGM is still perceived to be an intrinsic good and has a minimum value of z_3 otherwise) actors prepare to enter an initial coalition with other actors who are *conditionally ready* to abandon FGM. *'Organised Diffusion'* They also try to recruit zero-or-more (z_4) other actors in their household and social reference networks (at random) to join the initial coalition. If the recruiting actor's perceived intrinsic value of FGM is less than that of the recruitee, then the recruitee adopts a new (random uniform) belief in the value of FGM in the interval between theirs and the recruiting agents' belief. Recruitees respond (and prepare to join the initial coalition) according to the same probability rule as intervention participants. If they do respond, then they try to recruit zero-or-more (previously uninvolved) actors from *their* social and household reference groups to join the initial coalition. This continues until no more actors agree to join the initial coalition.

Update Enforcement Status	Actors who have joined the initial coalition update their status, such that pro-FGM enforcers stop enforcing the practice, and any actors who aren't anti-FGM enforcers become anti-FGM enforcers with probability $z_7 \in [0,1]$.
Connect coalition members	Each actor who is in the initial coalition adds a proportion (z_5) of actors in the coalition to their social reference group (at random).
Establish a stable (final) coalition	This is a recursive process: Actors in the initial coalition assess whether they would be willing to abandon FGM conditional on all other initial coalition members doing the same. Any actors for whom the answer is no, abandon the coalition. Remaining actors then repeat the same assessment. This continues until a final coalition stabilises, or the whole coalition collapses.

After implementing the simulated intervention, and in all other cases, the model steps through a series of discrete time-steps for as long as required (by default 10-years' worth of time-steps). In each time-step, actors in the population calculate an expected utility for practicing FGM and an expected utility for not doing so. These expected utilities depend on the intrinsic value that the actor attributes to FGM, and on social costs for practice or non-practice. Social costs are decomposed into:

1. Costs arising from social reference groups (immediate social consequences of miscoordination)
2. Costs arising from household reference groups (immediate social consequences of miscoordination)
3. Costs arising from decision-maker reference groups (anticipated future consequences for uncut girls based on the current rate of cutting among decision-makers that the agent is responsive to)

Immediate social consequences (derived from social and household reference groups) are further decomposed into those arising from:

- Implicit social pressures – derived from the authority-weighted proportion of actors engaged in FGM practice (or conversely FGM abandonment for social costs of *not* abandoning). Explicit social pressures (or norm enforcement) – derived from the authority-weighted proportion of actors *enforcing* FGM practice (or non-practice), with enforcement further weighted to create equivalence of total influence with passive enforcement scenarios (see section 7.11.2).

The relative importance of these different sources of social influence is controlled by scale parameters: $s_3, s_4, s_5 \in [0,1]$, such that social influence can be isolated in a particular type e.g. norm enforcement within the household, or spread across the different sources of influence.

If the expected utility for practicing FGM is greater than, or equal to, the utility for non-practice, then actors participate in FGM practice. If they are final decision maker about girls in their household then (with a certain probability that controls the rate of decision-making about girls, see 7.6, 7.12 and 8.4) they update their status to reflect that they decide to cut a girl in their household. Actors who participate and have the role of practice-enforcer, are also counted as enforcing the practice of FGM. Conversely, the situation is reversed if the utility for non-practice is greater: actors don't participate, may enforce non-practice (if applicable) and will update their cutting decision to non-cutting (if applicable and with a given probability).

Mover order is either simultaneous or sequential. If it is simultaneous, then all actors calculate their expected utilities, and subsequently, all actors make the decision to practice FGM or not. If it is sequential, then actors calculate utilities and act one at a time. The sequential move-order can either be random in each time step, or it can be in descending order of some characteristic (depending on the parameter a_4). Implemented possibilities are:

1. Random Sequential Move Order
2. Descending Order of Authority
3. Descending Order of Autonomy
4. Descending Order of Perceived Intrinsic Value of FGM

This process of utility calculation and decision-making is repeated once for every time-step until the simulation is stopped.

4. Design Concepts

The following table lays out descriptions of the 11 design concepts recommended by the ODD protocol.

Table 3: Design Concepts

DESIGN CONCEPT	DESCRIPTION
Basic Principles	The model is based on the framework provided by the social-norm of coordination account of FGM: that actors face social incentives to practice FGM if others do and will be willing to practice FGM if sufficient others do so. This conceptual framework has been shown to provide a superior account of the actor's underlying decision process than accounts focusing on informational influence or marriage *competition* in the literature (Chapter 4). This model extends the formal coordination model popular among policymakers (Mackie, 1996; UNICEF, 2007) by introducing: greater individual heterogeneity, different kinds of social reference group, complex local interaction structures (i.e. networks), variant decision-functions, additional effects from 'simulated interventions' and alternative temporal processes (e.g. staggered decisions about cutting relative to participation, and different move orders). Each of these extensions has been shown to potentially disrupt the predictions of the original model (Chapter 5 of the Thesis). The general model permits a space of possible combinations of these elaborations to be explored.
Emergence	The key emergent properties of the model are the rates of FGM participation that model arrives at given different starting conditions and/or simulated interventions. These emerge from, and can be disrupted by, a wide range of factors in ways that are not immediately obvious or centrally imposed by the model: including the details of actors decision-functions, the distribution of actor characteristics, the interactions structure of the population and features of the simulated intervention such as the actor-targeting strategy.
Adaptation	Actors continually adapt to the practice of FGM in their social, household and decision-maker reference groups (according to the priority given to these by global parameters in the simulation), since these determine social incentives which, in turn, affect their decisions.
Objectives	Actors continually try to maximise their expected utility by making a decision to practice or not practice FGM depending on the expected utility of both options.
Learning	Actors preference are updated by intervention participation, or in some cases, by social interactions surrounding an intervention (see *Process Overview* above). However, beyond this, the actors' preference and decision-functions are treated as stable and exogenous.

Prediction	Actors make implicit predictions when they respond to the rate of FGM practice on girls, based on their decision-maker reference group. The rate of cutting in this group reflects potential costs in the future (e.g. to marriageability) that girls may face if they are in a minority of their age-group who are uncut.
Sensing	Actors directly observe whether actors in their social and household reference groups participate in, and enforce, the practice of FGM (or conversely the abandonment of FGM). Actors are aware of whether decision-making actors in their decision-maker reference group cut their daughter the last time the decision arose. Actors do not respond directly to actors outside of these personal reference groups (although these *could* include the entire population).
Interaction	Interaction occurs primarily through actors observing the FGM participation/non-participation and practice or non-practice enforcement of one another in their reference groups and responding to the changing social incentives that this creates.
Stochasticity	The distribution of actor characteristics- authority, the intrinsic value of FGM, x & y coordinate positions in network space, and autonomy- is modelled stochastically using directly parameterised underlying marginal and joint probability distributions (based on the flexible beta-distribution family). Responsiveness to the intervention is also modelled stochastically since this is expected to be a variable process which depends on actor's attributes (i.e. their beliefs about FGM), but for which the explicit mechanism is not modelled.
Collectives	Households can be considered collectives. Households are small groups of actors whose household reference groups form a fully-connected network clique (all actors in the household are interconnected, but there are no external connections). One available intervention strategy is to target a proportion of households (rather than a proportion of actors) for intervention participation.
Observation	The primary data collected from the model are the modelling assumptions embedded in the global parameters (including related to intervention effects) and the corresponding resulting rate of FGM participation in the population after running the model (possibly after a simulated intervention) for a period of time.

5. Initialisation

[See *Process Overview and Scheduling*]

6. Input Data

See details of calibration in Chapter 7 of the thesis to which this model is attached.

7. Sub-models

7.1 Create Base Population

This is a simple sub-model which creates a population of agents. It also initialises the agent-attribute: `targeted?` as a Boolean variable which indicates (later in the simulation) whether the actor has been targeted for participation in an intervention, either by the observer or by other actors who are 'recruiting' after the initial intervention (i.e. organised diffusion, see 7.6). It also deals with some aesthetics of the output of the model. The following code implements this sub-model:

```
to create-base-population


  create-actors n-actors [
    [ code for aesthetics]
    set targeted? FALSE
  ]


end
```

7.2 Distribute Population Characteristics

7.2.1 Overall Implementation

The overall implementation of this sub-model depends on the following procedures:

1. `create-marginal-distributions` which creates lists of deviates from beta-distributions which act as approximations of the marginal distributions of the α_i, w_i and H_i characteristics of agents, with parameterised mean and variances.
2. `create-noisy-distributions` creates noise-disturbed versions of the marginal distribution of H_i which have parametrised correlations with the H_i distribution and which are used in the subsequent procedure (see also 8.2).
3. `assign-agent-attributes` assigns the H_i, α_i, w_i and y_i characteristics of agents in such a way that both the marginal distributions of these characteristics and their correlation with the H_i distribution is independently parameterised (see also 8.2).

7.2.2 Create Marginal Distributions

First, the procedure `create-marginal-distributions` creates the global list objects with 2000 deviates from the desired marginal distributions of agent characteristics. These marginal distributions are created using a version of the beta distribution family, which is parameterised in terms of its mean and variance (see 8.1).

Marginal distributions are created for the following agent-attributes

1. The intrinsic value of FGM (H_i)
2. Authority of actors (w_i)
3. Autonomy of actors (α_i)

This is implemented by the following Netlogo Code:

```
set marginal-distribution-H (n-values 2000 [x -> beta-draw mu_H sigma2_H] )
set marginal-distribution-authority sort (n-values 2000 [x -> beta-draw mu_w sigma2_w] )
set marginal-distribution-autonomy sort (n-values 2000  [x -> beta-draw mu_alpha sigma2_alpha] )
```

Note that this relies on the function `beta-draw` which takes as input a mean and variance for the beta distribution, and returns a single random deviate from that distribution. For each marginal distribution, there is a global parameter controlling the mean (`mu`) and the variance `sigma2`.

The beta-draw procedure is implemented by the following code:

```
to-report beta-draw [E_z V_z]


  ; Given a mean and variance value, this reporter acts as a random beta vari
able with corresponding mean and variance
  let x1 random-gamma (-( (E_z * ( (E_z ^ 2) - E_z + V_z)) / V_z)) 1
  let y1 random-gamma ( ((E_z - 1) * (V_z + (E_z ^ 2) - E_z))/(V_z) ) 1
  let z1 (x1 / (x1 + y1))
  report z1


end
```

To see why this produces a beta-distribution deviate with the desired mean and variance, consult section 8.1.

7.2.3 Create Noisy Distributions

After dealing with the marginal distributions of heterogeneous characteristics H_i, α_i and w_i, the model constructs objects necessary to manipulate the *joint* distribution of α_i, w_i and y_i (y_i is always uniformly distributed) with H_i. This section details with how this is implemented in the code, to see the mathematical reasoning which underlies this implementation, readers should consult section 8.2.

Construction of joint distributions begins with the sub-procedure `create-noisy-distributions`. This procedure is designed to create noise-disturbed versions of the marginal distribution of H_i (which was created by the `create-marginal-distributions` procedure). These distributions are disturbed by adding Gaussian noise (random normal distribution centred on zero) where that noise has a standard deviation determined analytically to produce a pre-specified Pearson product-moment correlation between the original distribution of H and the noise-disturbed distribution (see 8.2). As such, the procedure operates in two steps. First, the required standard deviation of the Gaussian noise is calculated for each of the joint distributions of interest ($p(H,\alpha), p(H,w)$, and $p(H,y)$). Second, a noise-disturbed version of the H_i marginal distribution is created for each of these joint distributions. This is implemented in the following Netlogo code:

```
 set desired-noise-sd-authority (sqrt(sigma2_H - (rho_Hw * sigma2_H))) / (rho_Hw)
 set desired-noise-sd-autonomy  (sqrt(sigma2_H - (rho_Halpha * sigma2_H))) / (rho_Halpha)
 set desired-noise-sd-Y (sqrt(sigma2_H - (rho_Hy * sigma2_H))) / (rho_Hy)

 set noisy-distribution-Q-for-authority map [x -> x + random-normal 0 desired-noise-sd-authority] marginal
-distribution-H
 set noisy-distribution-Q-for-autonomy  map [x -> x + random-normal 0 desired-noise-sd-autonomy] marginal-
distribution-H
 set noisy-distribution-Q-for-Y map [x -> x + random-normal 0 desired-noise-sd-Y] marginal-distribution-H
```

Derivation of the formula used to calculate the required standard deviation of the noise term in each case is provided in section 8.2.2.

7.2.4 Assign Agent Attributes

Having created three noise-disturbed versions of the H_i distribution in the model, one for each of the desired joint distributions: $p(H, \alpha, p(H, w)$, and $p(H, y)$, and each with a pre-specified Pearson produce moment correlation with the marginal distribution of H_i, we then assign the w_i, α_i and y_i attributes of agents in such a way that that they have a pre-specified non-linear correlation with H_i while maintaining their original marginal distributions. This is undertaken by the procedure `assign-agent-attributes`. To see the mathematical reasoning underlying this procedure, readers should consult section 8.2.

```
ask actors [

    set q_i one-of marginal-distribution-H ; Actors take a deviation from the marginal distribution of H_i (note that q_i
maps directly to a value of H_i later in the initialisation procedure)

    ; Each attribute is the inverse cumulative density function of the marginal distribution of the attribute, composed o
n the cumulative distribution of the noise-disturbed H_i distribution for that attribute, with a noise-disturbed deviate
from the H_i distribution as input...

    set alpha_i Inverse-CDF
                (CDF (q_i + random-normal 0 desired-noise-sd-autonomy) noisy-distribution-Q-for-autonomy) marginal-distri
bution-autonomy

    set w_i Inverse-CDF
                (CDF (q_i + random-normal 0 desired-noise-sd-authority) noisy-distribution-Q-for-authority) marginal-dist
ribution-authority
```

```
    set y_i CDF (q_i + random-normal 0 desired-noise-sd-y) noisy-distribution-Q-for-y ; The inverse CDF of a uniform dist
ribution is simply the CDF (y_i is always random-uniform)

    set x_i random-float 1
  ]
```

This procedure relies on functions to approximate the cumulative and inverse cumulative density functions numerically (since these are not known analytically). The function `CDF` takes as input an empirical marginal distribution `dist` and a deviate from that distribution `q`. It returns the CDF of that distribution evaluated at `q`:

```
to-report CDF [q dist]
  report (length filter [i -> leq i q] dist) / 2000
end


to-report leq [i x]
  report i <= x ; returns a value TRUE or FALSE
end
```

The function `Inverse-CDF` takes as input the output of a CDF function `cdf_x` and a marginal distribution `dist`. It returns (a numerical approximation of) of the inverse cumulative distribution function of `dist` evaluated at `cdf_x`:

```
to-report Inverse-CDF [cdf_x dist]
  report item (cdf_x * 1999) dist ; note that the constant 1999 is based on t
he size of the simulated marginal distributions (2000 deviates)
end
```

The key point that readers should understand is that these procedures (within `distribute-population-characteristic`) effectively parameterised the marginal distributions of the α_i (autonomy: $\mu_\alpha, \sigma_\alpha^2$) w_i (authority: μ_w, σ_w^2) and H_i (perceived intrinsic value of FGM: μ_H, σ_H^2) attributes of agents, as well as the strength of the correlation of the joint distributions of H_i with attributes: α_i (intrinsic value and autonomy correlation: $\rho_{H\alpha}$), w_i (intrinsic value and authority correlation: ρ_{Hw}), and y_i (intrinsic value and y-coordination position correlation: ρ_{Hy}).

The `assign-agent-attributes` procedure further assigns the `x_i` attribute of actors, which is a random uniform deviate in the interval $[0,1]$.

Readers should note that at this stage in the initialization of the model, the attributes: α_i, w_i, q_i (which is the random deviate that determines H_i) and y_i are all in the interval $[0,1]$. These attributes are scaled as required during the `setup-initial-behaviour` procedure later in the simulation.

7.3 Set Spatial Locations

This is a simple sub-model which maps the y_i and x_i attributes of actors to their location in a 2D Cartesian coordinate space (called the *network space*) which wraps on its horizontal axis. Dimensions are arbitrary. This is implemented in the following Netlogo code:

```
to set-spatial-location


    ask actors [
    set xcor x_i * (world-width - 1)
    set ycor y_i * (world-height - 1)
  ]


end
```

It is important for readers to note that if ρ_{Hy} has a positive value, such that y_i and H_i are correlated in the simulation, then actors' position in the vertical dimension will be correlated with their perceived intrinsic value of FGM. This, in turn, will affect the level of homophily within their social and household reference groups (see also: 7.4).

7.4 Form Reference Networks

7.4.1 Overall Implementation

This sub-model depends on three separate procedures, one for each of the reference groups of actors:

1. `create-household-reference-networks` - This procedure connects actors in random network cliques within partitions of the network space. All actors within each clique form one-another's household reference group ($\{\text{household}\}_i$). Within these households, a small (random) number of actors are assigned the role of decision-maker about the FGM status of some of the girls in that household ($m_i = 1$).
2. `create-decisionmaker-reference-network` - This procedure connects actors to all other actors who are decision-makers and are within a fixed radius of them in network space. These connections define actors' decision-maker reference groups ($\{\text{decision-maker}\}_i$).
3. `create-social-reference-network` - This procedure connects actors to all other actors that are not part of their household and are within a fixed radius of them in network space. These connections define the actors' social reference groups ($\{\text{social-ref}\}_i$).

7.4.2 Create Household Reference Networks

This procedure creates households, which are small network cliques within the network space, where each actor in the household is in the other actor's household reference group, and there are no external connections (of the household reference group type). The algorithm is flexible with respect to whether network cliques are formed randomly in the population (implying no relationship between social and household networks), or whether cliques are formed within localised partitions of the network space - implying that there is a close correspondence between social and household networks. It is worth noting that in the former case, where household cliques span the entire network space, these cliques effectively bridge disparate parts of the social reference network. In this case, there will be no homophily of preferences (or other characteristics) of householders. Conversely, when household cliques are localised within partitions of the network space, they only bridge relatively close parts of the social network, and homophily in the social network will translate into homophily of the household networks. The degree of localization of household cliques is controlled by the parameter η_1, which determines the number of (equally spaced) horizontal partitions of the network space in which cliques are randomly formed. Pseudo-code for the algorithm that forms the household cliques is as follows:

```
for each of eta_1 equally spaced horizontal partitions of network space:
    while there are agents without a household in that space:
        ask up to random Poisson mu_hsize actors in the partition:
            form a network clique and become one another's household reference group
```

Where `mu_hsize` is the average number of actors in each household clique.

After forming each of the households, a small number of actors in each household (typically 2) are assigned the role of decision-maker about (some) girls in that household ($m_i = 1$). This is a pseudo-code representation of this process:

```
for each household in the population:
    ask 1 + random-Poisson 1 householders:
        set decision-maker (m_i) 1
```

7.4.3 Create Decision-Maker Reference Networks

This procedure creates decision-maker reference groups for agents by connecting agents to decision-maker agents within a fixed Euclidean distance of them in network space. This distance is parameterised to be a scaled version of the distance used for the social reference group (which follows a similar procedure), where the scaling factor: $s_8 \in [1,3]$ is always greater than or equal to 1. We will call the distance used in constructing the social reference network the *social-reach* of actors, and the distance used in constructing the decision-maker reference network the *decision-maker-reach*, noting that *decision-maker-reach* is equal to *social-reach* multiplied by s_8.

Therefore, to implement the decision-maker reference network, we need to first calculate the *social-reach*. Rather than being parameterised directly, this distance is parameterised in terms of the desired average connectivity of agents in the social reference group. This is done to ensure that the connectivity of the social reference network is invariant to parameters controlling the size of the population. As such, the distance is determined by a parameter $\mu_{R_{social}}$ representing the expected average size of actors' social reference group. The calculation is implemented in the Netlogo code as:

```
set social-reach-of-actors social-reach-formula n-actors mu_Rsocial

to-report social-reach-formula [n c_i]

  let A_t (world-width * world-height)
  report sqrt ((A_t * c_i) / (pi * (n - 1)))

end
```

Note that the `social-reach-formula` function takes as input the number of actors in the population `n-actors` and the expected average size of actor's social reference networks $\mu_{R_{social}}$, and then returns the *social-reach* value that will produce this result (on average). To see the derivation of this formula, readers should consult section 8.3.

Having calculated *social-reach*, we derive *'decision-maker'-reach* and ask actors to add all other decision-maker actors within a distance of *'decision-maker'-reach* in network space to their decision-maker reference group:

```
ask actors [
    set R_decision_maker other actors with [distance myself <= (social-reach-of-actors * s_8) and m_i? = 1
]
    set R_decision_maker_N count R_decision_maker
  ]
```

7.4.4 Create Social Reference Networks

The operation of this procedure is very similar to `create-decisionmaker-reference-networks`. Actors add all other actors within a fixed distance in network-space to their social reference group ($\{\text{social-ref}\}_i$) provided that those actors are not already part of their household reference groups. This is implemented by the following Netlogo code:

```
ask actors [
    let my-household R_household
    set R_social other actors with [distance myself <= social-reach-of-actors
and not member? self my-household]
    set R_social_N count R_social
  ]
```

7.5 Setup Initial Behaviours

7.5.1 Overall Implementation

This sub-model acts as a single procedure. However, it has a number of conceptually separate components:

1. Scale Agent Attributes - Continuous attributes are scaled to an appropriate interval
2. Assign Enforcement Roles - A certain proportion of agents' roles as enforcers of practice or non-practice
3. Initialise Behaviors - A certain proportion of agents are set to initially practice FGM

7.5.2 Scale Agent Attributes

In previous procedures (7.2), agents' q_i attributes were set, where this is a random deviate that determines their H_i attribute (their perceived intrinsic value of FGM). In this procedure their H_i attribute is formally assigned. This can be seen as scaling their q_i attribute to an interval $[-\widehat{M}, \widehat{M} \cdot s_2]$, where s_2 is a scaling parameter in $[0,1]$ that determines the size of the maximum perceived value of FGM, relative to its minimum perceived value (i.e. maximum perceived cost). This is achieved through defining H_i as follows:

$$H_i = -\widehat{M} + q_i \cdot (\widehat{M} + [\widehat{M} \cdot s_2])$$

This is implemented using the following NetLogo code:

```
ask actors [
    set H_i Q-to-H q_i
    ]

to-report Q-to-H [Q]
  report (-1 * M_hat) + (Q * (M_hat + (M_hat * s_2)))
end
```

Readers should note the following properties of this scaling procedure:

1. If s_2 is 0, then no actors will view FGM as having intrinsic value
2. IF s_2 is 1 and the PDF of q_i is symmetric, then perceived intrinsic value of FGM in the agent population (H_i) will be symmetrically distributed around the value of $H_i = 0$.

The other agent attribute which is scaled by this procedure is α_i (actor's autonomy). This is scaled from the $[0,1]$ interval to an $[-1,1]$ interval as follows:

```
ask actors [
    set alpha_i (-1 + (2 * alpha_i))
  ]
```

7.5.3 Assign Enforcement Roles

In this procedure, a certain proportion of agents (s_6) who view FGM as an intrinsic good ($H \geq 0$) are assigned the role of enforcer of the practice ($\beta_{1i} = 1$), which means that if they practice FGM, they will convey active social influence on others to do the same (see 7.11). Conversely, a certain proportion of agents (s_7) who view FGM as an intrinsic ill ($H < 0$) will be assigned the role of enforcers of abandonment of the practice ($\beta_{2i} = 1$, which means that if they do not practice FGM they will convey active social influence to others to abandon the practice (see 7.11).

In each case (enforcing practice or non-practice) agents are assigned to these roles according to a rule for preferentially selection agents. Three possible rules are supported in the model:

1. "random" - meaning that agents are selected uniformly at random
2. "zealots" - agents that are most extreme in their beliefs about the intrinsic value of FGM (either positive or negative) are preferentially selected as enforcers or practice or non-practice (respectively).
3. "high-authority" - agents with the highest authority in the population w_i are preferentially selected as enforcing FGM

These options are stored in parameters a_1 and a_2 respectively for FGM enforcement and abandonment enforcement. Assignment of these roles is implemented by the following Netlogo code:

```
   if a_1 = "random" [if any? supporters [ask n-of (s_6 * (count supporters)) supporters [set beta_1i? 1]
] ]
  if a_1 = "zealots" [if any? supporters [ask max-n-of (s_6 * (count supporters)) supporters [H_i] [set be
ta_1i? 1] ] ]
  if a_1 = "high-authority" [if any? supporters [ask max-n-of (s_6 * (count supporters)) supporters [w_i]
[set beta_1i? 1] ] ]

  if a_2 = "random" [if any? opponents [ask n-of (s_7 * (count opponents)) opponents [set beta_2i? 1] ] ]
  if a_2 = "zealots" [if any? opponents [ask min-n-of (s_7 * (count opponents)) opponents [H_i] [set beta_
```

```
2i? 1] ] ]
  if a_2 = "high-authority" [if any? opponents [ask max-n-of (s_7 * (count opponents)) opponents [w_i] [se
t beta_2i? 1] ] ]
```

Where `supporters` is a list of agents for whom H_i is greater than or equal to zero, and `opponents` is a list of agents for whom H_i is less than zero. Readers should note that if there are no agents who support FGM, the first section of code does nothing.

7.5.4 Initialise Behaviours

This function assigns a certain proportion of actors `initial-participation-rate` as practising FGM at the beginning of the simulation. There are two possible options for this assignment, controlled by the parameter a_3. Either this assignment is random $a_3 = \text{random}$, or actors are selected in descending order of H_i (such that supporters of FGM are preferentially selected as practising FGM at the start of the simulation, $a_3 = \text{intrinsic-value-first}$). In either case, selected actors set their FGM participation to 1 ($d_i = 1$) and if they are a decision-maker about the FGM status of (some) girls in their household ($m_i = 1$), they set their cutting decision to 1 ($\theta_i = 1$).

7.6 Setup Cutting Decision Process

In the model, decision-makers decisions about whether to cut their daughters are determined by their decisions about whether to participate in FGM activities. There is no separate utility calculation from the perspective of the decision-makers (although the cutting decisions of decision-makers directly affect the utility calculations of agents for whom those decision-makers are part of a decision-maker reference group, see 7.11 and 7.4.3).

However, participation (updating d_i) and cutting decisions (updating θ_i) take place on different time scales. While all agents have the opportunity to participate in FGM in each time-step of the model, decision-makers only have the opportunity to update their cutting decisions with a fixed probability that is determined analytically to reflect the overall rate of cutting decisions taking place in the community per year. The derivation of this probability is given in section 8.5. Here we deal with the mechanics of calculation within the `setup-cutting-decision-process` procedure.

In essence, this procedure calculates the number of girl children in the community, based on the number of households (and the number of children per household, see 8.5). It then calculates the number of decisions to be taken per-year, given that all girls must have a final decision made by age 15. Then based on (a) the number of decisions to be taken each year, (b) the number of simulated time steps per year (controlled by the parameter $\eta_2 \in 1,2,...,11,12$) and (c) the number of decision-makers among whom these decisions are distributed, the probability that decision-makers update their decision in each time step is determined analytically. This is implemented in the following Netlogo code:

```
to setup-cutting-decision-process

  set est-number-of-children (length family-list) * 0.84 * 2.25 ; This is based on the average
proportion of houses with children, and the average number of children per house
  set est-decisions-per-year (est-number-of-children / 15) ; Each year at least one-fifteenth
of children must be cut to account for the rate of arriving children (implicit)
  let number-dms count actors with [m_i? = 1]
  set p-decide-per-step (est-decisions-per-year) / (eta_2 * number-dms) ; Ensures that on aver
age, 'est-decisions-per-year' are made regarding the cutting status of girls

end
```

When the model steps through a time-step (see 7.12), actors who are decision-makers about (some) girls in their household, will update their cutting decision (θ_i) to reflect their last participation decision (d_i) with probability `p-decide-per-step`.

7.7 Target Actors for Intervention

This is the first procedure in the implementation of a simulated intervention. Simulated interventions happen before the model steps through time, and we assume that the model is initialised such that all actors are practicing FGM.

This first procedure deals with which agents are initially targeted to participate in the intervention and the initial educational effects of the intervention on those agents. A wide range of options are supported by the model and the rule used is controlled by a categorical variable: z_6. A proportion z_1 of agents can be targeted at random, or agents can be targeted as follows:

- Support FGM (H_i descending order)
- Oppose FGM (H_i ascending order)
- Autonomy (α_i descending order)
- Authority (w_i descending order)
- Pro-FGM enforcement (${\beta_1}_i = 1$ selected)
- Anti-FGM enforcement (${\beta_2}_i = 1$ selected)
- Social connectivity ($\{\text{social-ref}\}_i$ size, descending order)
- Household Size ($\{\text{household}\}_i$ size, descending order)
- Decision-makers ($m_1 = 1$ selected first)
- Localised (x_i descending order)
- By Household (a proportion z_1 of households are targeted, and all actors in those households participate in the intervention).

After a proportion z_1 of actors are targeted to participate in the intervention, these actors set their status `targeted?` to `TRUE`. These participating actors then set their belief about the value of FGM according to the following procedure (representing the educational component of the intervention):

```
ask actors with [targeted? = TRUE] [set H_i intervention-H-change q_i]

to-report intervention-H-change [q]
 report Q-to-H ((1 - z_2) * q)
end
```

```
to-report Q-to-H [Q]
  report (-1 * M_hat) + (Q * (M_hat + (M_hat * s_2)))
end
```

This effectively scales down actors' perception of the value of FGM by a factor of $1 - z_2$ where z_2 is a global parameter indicating the strength of the educational component of the intervention. When z_2 is 1, all participating actors will set their belief in the intrinsic value of FGM to the lowest possible value: $H_i = -\widehat{M}$. When z_2 is 0, participating actors' H_i preference will not change. The relevant formulas are:

Agent's perception of the value of FGM *before* the educational effect of the intervention is defined:

$$H_i = -\widehat{M} + (q \cdot \left[\widehat{M} + \left(\widehat{M} \cdot s_2\right)\right])$$

After the educational effect of the intervention, it is defined:

$$H_i = -\widehat{M} + ([(1 - z_2) \cdot q] \cdot \left[\widehat{M} + \left(\widehat{M} \cdot s_2\right)\right])$$

7.8 Implement Intervention

Actors who are targeted by the intervention all receive the educational component of the intervention. However, they then make a decision about whether to respond to the intervention by spreading the educational component (optional on parameter z_4) to their reference groups, and trying to form a coalition of abandonment with other practitioners. The mechanism of this decision isn't modelled directly. Instead, I assume:

1. Actors who still believe that FGM is an intrinsic benefit will not try to arrange a coalition to abandon it (since this is not their preferred outcome)
2. The likelihood of actors forming an initial coalition of others willing to abandon FGM is a function of the strength of their perception of the cost of FGM.

Actors' decision to respond to the intervention is then modelled using the following rule:

```
if H_i < 0 and random-float 1 < (1 - (( -1 * ((1 - z_3) /  M_hat)) * H_i)) [

    respond to the intervention
```

This essentially defines actors' probability of responding to the intervention as 0 if they prefer FGM ($H > 0$) and an increasing linear function of their perception of the intrinsic *cost* of FGM otherwise:

$$\Pr(join\ initial\ coalition)_i = \begin{cases} -\left(\frac{1 - z_3}{\widehat{M}} \cdot H_i\right) + z_3\ if\ H_i < 0 \\ 0\ if\ H_i \geq 0 \end{cases}$$

This linear function is controlled by a parameter z_3. It is easiest to think of this as the minimum probability that actors who oppose FGM will respond to the intervention - since it represents the probability of response when H_i approaches 0 from a negative direction. The probability of responding to the intervention is always 1 when actors attribute the maximum social cost to FGM: $H_i = -\widehat{M}$. Also, if z_3 is 1, then actors always respond to the intervention.

Actors who respond to the intervention set their `coalition-ready?` status to `TRUE` to reflect that they are prepared to enter a coalition of actors abandoning FGM. They also try to influence and recruit others to join the coalition. This represents the much-discussed intervention feature 'organised diffusion' (see Chapter 5 and UNICEF, 2007).

Actors select z_4 random other actors from their social and household reference groups (provided those actors were not part of the intervention and no one else has tried to recruit them) and 'recruit' them (if z_4 is set to 0, then this procedure has no effect on the simulation). Then, if any of these recruitees has a higher H_i attribute than the recruiting actor, the recruitee sets their own H_i value to a random uniform location in the interval between the recruiting and recruitee actors' H_i attribute. This represents the recruiting actor influencing the beliefs of the recruitee (i.e. spreading the educational component of the intervention).

Subsequently, recruitees decide whether they want to respond to the intervention by preparing to enter an abandonment coalition, and influencing others (etc.). Their decision follows the same probabilistic rule as the original intervention participants (see above).

As such, the `implement-intervention` procedure continues recursively: with actors responding or not, and recruiting others who then respond or don't, and so on until no more actors respond. This is implemented by the following Netlogo code:

```
to implement-intervention

  if H_i < 0 and random-float 1 < (1 - (( -1 * ((1 - z_3) /  M_hat)) * H_i)) [

    set coalition-ready? TRUE

    let recruitment-agentset (turtle-set R_social R_household)

    if count recruitment-agentset with [targeted? = FALSE] > 0 [
      let my-H_i H_i
      ask up-to-n-of z_4 recruitment-agentset with [targeted? = FALSE] [
        if H_i > my-H_i [set H_i (my-H_i + random-float (H_i - my-H_i) ) ]
        set targeted? TRUE
        implement-intervention
      ]
    ]
  ]

end
```

7.8b Update Enforcement Status

Actors in the initial coalition update their status with respect to the enforcement of FGM. Actors who are pro-FGM enforcers, drop this status $\beta_{1_i} \rightarrow 0$. Actors then become anti-FGM enforcers ($\beta_{2_i} \rightarrow 1$) with probability z_7.

7.9 Connect Initial Coalition Members

The model allows for the possibility that participation in the intervention changes the social reference network structure of the population. This is controlled by the parameter z_5. If z_5 is 0 then there is no effect on the simulation. Otherwise, actors add a proportion z_5 of those who have responded to the intervention (i.e. are ready to join an abandonment coalition) and who are not part of their household reference group, to their social reference group:

```
let coalition-group actors with [coalition-ready? = TRUE]
   let coalition_N count coalition-group
   ask coalition-group [
    let my-household household-label
    set R_social (turtle-set R_social (up-to-n-of (z_5 * coalition_N) (other
coalition-group with [household-label != my-household]) ))
    set R_social_N count R_social
    set R_social_weight sum [w_i] of R_social
  ]
```

7.10 Establish Coalition

This procedure deals with the formation of a stable final coalition among initial coalition members. The rule for stable coalition formation is a simple one: it is stable if all actors in the coalition prefer to abandon FGM, conditional on all others in the coalition doing so. This is modelled as a recursive process. Initial coalition group members assess whether they would prefer to abandon FGM, conditional on all others in the coalition doing the same. If not, they leave the initial coalition. Remaining members repeat the assessment (based on the remaining number of coalition members). This is repeated till a final coalition stabilises, or collapses entirely.

To find a stable coalition (if it exists at all), the following algorithm is used (pseudo-code):

```
 Let C be the set of actors initially willing to enter a coalition
  1. Members of C visibly stop any FGM activity
  2. Members calculate whether, under these conditions, they would prefer to
abandon FGM

  3. While there are some agents in the coalition who would prefer to keep pr
acticing FGM:
     -> Agents who still prefer (bc. of social incentives) to practice FGM, r
evert to practicing
     -> Agents who still prefer to practice FGM leave the coalition
     -> Remaining members of the coalition calculate whether they still prefe
r to abandon FGM (given the reduced coalition size)

   4. The remaining members of C represent a stable coalition of actors aband
oning FGM.
```

This algorithm finds the stable abandonment coalition if it exists, and these actors visibly abandon FGM activities ($d_i = 0$ and $\theta_i = 0$).

7.11 Calculate Expected Utilities

7.11.1 The decision-functions of agents

The decision-function of agents was implemented as follows:

$$U(abandon)_i = p_{pro_i}^{V^{(-\delta_1+\alpha_i)}} \cdot -\widehat{M}$$

$$U(practice)_i = H_i - (s_1 \cdot p_{anti_i}^{V^{(-\delta_2+\alpha_i)}} \cdot \widehat{M})$$

Where H_i is defined as:

$$H_i = -\widehat{M} + (q \cdot [\widehat{M} + (\widehat{M} \cdot s_2)])$$

Note here that H_i is defined as the intrinsic 'value' of FGM to the agent (rather than the intrinsic cost), which can be positive (FGM is viewed as intrinsically beneficial) or negative (FGM is viewed as intrinsically costly). Since H_i can be positive or negative in the general model, I adopt the term 'FGM supporters' to refer to actors who value FGM ($H_i \geq 0$) and the term 'FGM opposers' to refer to actors who view FGM as costly ($H_i < 0$)[1].

These functions have the following properties:

- H_i can vary between $-\widehat{M}$ and $\widehat{M} \cdot s_2$. As such, the $s_2 \in [0,1]$ parameter controls the maximum perceived positive value of FGM in the population.
- The maximum social cost actors pay to unilaterally practice FGM varies from 0 to $\widehat{M}$ and is controlled by the s_1 parameter.
- The relationship between social pressure (p_{anti_i}/p_{pro_i}) and social costs ($\widehat{M}$) is controlled by parameters V (≥ 1), δ_1 (or δ_2, both in the interval $[-1,1]$) and α_i ($\in [-1,1]$, representing the autonomy of individual actor i). V controls the overall nonlinearity of the relation between social pressure and social costs. δ_1 and δ_2 allow for global control of the relation between social pressures and social costs for abandoning or practicing FGM (respectively). As δ_1 increases, for example, the costs to abandon FGM 'scale' faster with social pressure, making

[1] This replaces the use of the term 'willing agents' and 'reluctant agents' in Chapter 5.

it harder for actors to abandon the practice (and vice-versa for δ_2). δ_1 and δ_2 allow the model to accommodate the possibility that the relationship between social pressure and social costs varies *overall* for those practicing versus abandoning FGM. The effect of these parameters is eliminated if they are set to 0. The α_i component allows variation in the social-pressure social-cost relationship at the individual level, with costs scaling faster for less autonomous agents (lower α_i).

The utility-functions of agents in the standard model are a special case of this more general formulation, in which $V = 1$, $s_1 = 0$, and $s_2 = 0$. Other previously seen variations on the decision process can be achieved by appropriate manipulations of these parameters.

In the above formulation, I don't define p_{anti_i} and p_{pro_i}, beyond the obvious: that they represent pro-FGM and anti-FGM activity by others. However, I assume that they are both bounded between 0 and 1. We can turn now to their (general) definition in the model.

7.11.2 Different Sources of Sources of Social Influence

I assume that influence from cutting decision-makers and influence from households/social reference groups are broadly distinguishable. The latter is primarily a source of immediate normative social pressure, while the former is about the future state of the marriage-market/social situation of girls if they are not cut. I controlled the relative influence of each with a parameter $s_4 \in [0,1]$. I then further distinguished between social influence from within households, versus the wider social reference group. I controlled the relative influence of each with parameter $s_3 \in [0,1]$. A pseudo-code representation of this would be:

$$[\text{total social influence}]_i = s_4 \cdot [\text{decision-maker influence}]_i + (1 - s_4) \cdot [\text{normative social influence}]_i$$

Where:

$$[\text{normative social influence}]_i = s_3 \cdot [\text{household influence}]_i + (1 - s_3) \cdot [\text{social reference group influence}]_i$$

Given the characterization of decision-maker influence (above), I defined pro-FGM influence from this source as the proportion of decision-makers who cut their daughters the last time the decision arose (see below), and anti-FGM influence from this source as the proportion who *didn't* cut their daughters last time the decision-arose.

Other sources of social influence (i.e. normative social influence) incorporated heterogeneous weights, representing the authority of individual actors. These other sources of social influence were also divided into explicit social influence (i.e. FGM practice by pro-FGM norm enforcers or FGM abandonment by anti-FGM norm enforcers) and implicit social influence (i.e. FGM practice in general), controlled by parameter $s_5 \in [0,1]$, for example:

$$\text{[social reference group influence]} = s_5 \cdot \text{[explicit 'norm enforcement.']} + (1 - s_5) \cdot \text{[implicit social influence]}$$

In the case of explicit social influence, norm-enforcers were re-weighted such that the total explicit social influence from all pro/anti-FGM enforcers (whichever group was larger) was equal to the total implicit influence of all actors (see below).

Formal definitions of all of the components of p_{pro_i} (total pro-FGM social pressure facing actor i) were as follows[2].

Let A_i be equal to the total implicit pro-FGM influence from the social reference group ($\{\text{social-ref}\}_i$) of actor i:

$$A_i = \frac{\sum_{j \in \{\text{social-ref}\}_i} w_j \cdot d_j}{\sum_{j \in \{\text{social-ref}\}_i} w_j}, i \neq j$$

Where $w_i \in [0,1]$ is the authority weight of actor j and d_j is a decision-indicator which is 1 when practicing FGM and 0 otherwise.

Let B_i be equal to the total *explicit* pro-FGM influence from the social reference group of actor i:

$$B_i = \frac{\sum_{j \in \{\text{social-ref}\}_i} w_j \cdot d_j \cdot \beta_{1j} \cdot w_2}{\sum_{j \in \{\text{social-ref}\}_i} w_j}, i \neq j$$

Where β_{1j} is a dummy indicator that is 1 if actor j is a pro-FGM norm enforcer and 0 otherwise. w_2 is a weighting coefficient chosen such that, the total influence 'weight' of all pro/anti-FGM

[2] Readers can substitute in $(1-d)$ for d, $(1-\theta)$ for θ and β_2 (indicating an anti-FGM norm enforcers) for β_1 for the full definition of p_{anti_i}.

enforcers in the community (whichever group is larger), is equal to the total influence 'weight' of all actors[3] (see below).

Let C_i be the total *implicit* social influence from the *household* and let D_i be the total *explicit* social influence from the *household*. These are defined in the same way as A_i and B_i (respectively), except that $\{social - ref\}_i$ is replaced with $\{household\}_i$ which is the set of other actors in the household of actor i.

Let E_i be the total social influence from the set of decision-makers that actor i is responsive to (see decision-maker reference group below): $\{\text{decision-makers}\}_i$. This is defined as:

$$E_i = \frac{\sum_{j \in \{\text{decision-makers}\}_i} \theta_j}{\sum_{j \in \{\text{decision-makers}\}_i} 1}, i \neq j$$

Where θ_j is a dummy variable indicating that decision-maker agent j cut a girl in their household the last time the decision arose.

We can then define p_{pro_i} as follows:

$$p_{pro_i} = (s_4 \cdot E_i) + (1 - s_4) \cdot ([1 - s_3] \cdot [s_5 \cdot B_i + (1 - s_5) \cdot A_i] + s_3 \cdot [s_5 \cdot D_i + (1 - s_5) \cdot C_i])$$

As noted above, I also defined a weighting coefficient, such that the total weighted influence of the largest group of enforcers (pro or anti-FGM, whichever was larger) was equivalent, on average, to the total *implicit* influence of all actors. This maintains the conformist properties of the simulation, and ensures, on average, that social influence is not biased in favour of implicit or explicit social influence (instead, this is explicitly controlled by parameter s_5.

The weighting coefficient for norm enforcement (w_2) was defined as:

$$w_2 = \frac{n}{n_{enforcers}}$$

---

[3] In the event that, due to stochastic effects, the total norm enforcement influence exceeds 1 (e.g. because the actor is connected to an unusual number of enforcers in the network), the influence is capped at 1 by the simulation.

Where n is the number of agents, and $n_{enforcers}$ is the number of pro-FGM or anti-FGM enforcers (whichever is larger), called 'the largest enforcement group'. Using w_2 as a weighting coefficient ensures that, on average, the total influence of the largest enforcement group is equal to the total influence of all actors (under implicit enforcement).

Under implicit enforcement, the average (i.e. the *expected*, note the $E[\cdot]$ operator) total influence of all actors is:

$$E\left[\sum_{i=1}^{n} w_i\right] = \sum_{i=1}^{n} w_i$$

Where w_i is the weight of actor i.

Under explicit enforcement, the total influence of the largest enforcement group is:

$$E\left[\sum_{i=1}^{n} w_i \cdot w_2 \cdot x\right]$$

Where x is a random dummy variable (0 or 1) that indicates that actor i is in the largest enforcement group. The expectation of x is $\frac{n_{enforcers}}{n}$. Extracting w_2 and x from the summation we find:

$$E\left[\sum_{i=1}^{n} w_i \cdot w_2 \cdot x\right] = E\left[w_2 \cdot x \cdot \sum_{i=1}^{n} w_i\right] = w_2 \cdot E[x] \cdot E\left[\sum_{i=1}^{n} w_i\right]$$

Substituting the definition of w_2 and evaluating $E[x]$ we find:

$$w_2 \cdot E[x] \cdot E\left[\sum_{i=1}^{n} w_i\right] = \frac{n}{n_{enforcers}} \cdot \frac{n_{enforcers}}{n} \cdot E\left[\sum_{i=1}^{n} w_i\right]$$

This then reduces to:

$$1 \cdot \sum_{i=1}^{n} w_i$$

7.12 Make Decision

If $U(practice)_i$ is greater than or equal to $U(abandon)_i$ then agents will practice in FGM; otherwise, they will not. Actors assigned the role of 'decision-maker' ($m_i = 1$) will update their more recent decision about the cutting status of a girl in their household (θ_i) to reflect their practice decision, with probability: $\frac{l}{x \cdot m}$ (see 8.4).

8. Model Discussion, Sub-Model Analysis & Theorems

8.1 Generalizing Over Continuous Marginal Distributions using the Beta Family

The model generalises over the marginal distributions of the continuous attributes of actors (H_i, α_i, w_i) in the modelled population as follows. The beta distribution was used. It allows a wide variety of forms, including unimodal, bimodal, left-skewed and right-skewed. However, it presents further challenges in that it is difficult to generalise over the space of beta distributions because its standard parametrization is unbounded (i.e. parameters can be infinitely large).

The following outlines the beta family of distributions and demonstrates how it can be re-parameterised in terms of its mean and variance. These are bounded values which allow a full exploration of the space of distributions. Furthermore, I show how this distribution can be decomposed into gamma distributions which allowed implementation in the Netlogo programming environment used to implement the model.

The beta distribution is a family of probability distributions for which all values outside of 0 and 1 have a zero density. It has PDF:

$$f(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)}$$

Where:

$$B(\alpha,\beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$$

and Γ represents the *gamma function*.

The beta-distribution family spans a wide variety of qualitatively distinct shapes, so can be used to approximate a range of distributions of interest for any finite range of values. The beta distribution is determined by two parameters: α and β. It also has some known dispersive properties in relation to these parameters. Specifically, if X is a beta-distributed random variable, then:

$$E[X] = \frac{\alpha}{\alpha + \beta} = E_X$$

Furthermore:

$$V[X] = \frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)} = V_X$$

The main practical barrier to using the beta distribution in a simulation context was that the values of α and β are *unbounded*; they can be any positive real number. As such, it was difficult to define a finite 'space' of beta-distributions which can be explored in the simulation. The *solution* to this issue was to define the beta distribution *in terms of* E_X and V_X instead. These values have a more meaningful interpretation (centre and spread), and they are bounded (such that the whole space of distributions can be explored systematically).

We know that E_X is bounded in the interval $[0,1]$, we also know from the *Popoviciu inequality* that the maximum variance of a bounded probability distribution is:

$$\frac{1}{4}(M - m)^2$$

Where M is the upper bound and m is the lower bound. Since these values are 1 and 0 for the beta distribution, this simplifies to a maximum variance of $\frac{1}{4}$ for the beta distribution.

To define the beta distribution in terms of its variance and mean, we simply apply methods for simultaneous equations:

$$E_X = \frac{\alpha}{\alpha + \beta}$$

$$V_X = \frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)}$$

First, we define α in terms of β and E_x. Note that here we restrict the solution to the case in which E_X is not equal to 1 (although, of course, it can be arbitrarily close to 1). β is always greater than zero.

$$\alpha = -\frac{\beta E_X}{(E_x - 1)}$$

Then we substitute the definition of α into the variance equation, and solve for β:

$$V_X = \frac{-\frac{\beta E_X}{(E_x - 1)}\beta}{(-\frac{\beta E_X}{(E_x - 1)} + \beta)^2(-\frac{\beta E_X}{(E_x - 1)} + \beta + 1)}$$

$$\Rightarrow \beta = \frac{(E_X - 1)(E_X^2 - E_X + V_X)}{V_X}$$

Substituting this definition of β back into the definition of α, we are left with two definitions of these parameters purely in terms of the desired variance and mean of the distribution:

$$\alpha = -\frac{\frac{(E_X - 1)(E_X^2 - E_X + V_X)}{V_X}E_X}{(E_x - 1)} = -\frac{E_X(E_X^2 - E_X + V_X)}{V_X}$$

$$\beta = \frac{(E_X - 1)(V_X + E_X^2 - E_X)}{V_X}$$

We could then redefine the beta-distribution family in terms of these values, giving them a more intuitive specification. However, in the use-case of interest here, we are using built-in functionality from Netlogo to create a beta distribution. Specifically, we are interested in defining our beta distribution in terms of the *Gamma* distribution (which is supported directly in Netlogo). Here we rely on the following relation between the *Gamma* and *Beta* distributions: if X and Y are independent random variables, where $X \sim \Gamma(\alpha, \theta)$ and $Y \sim \Gamma(\beta, \theta)$, and where α and β are the corresponding parameters of the beta distribution, then the random variable:

$$Z = \frac{X}{X + Y} \Rightarrow Z \sim Beta(\alpha, \beta)$$

This holds irrespective of θ.

Based on the above, we can define the following random variable in terms of the desired mean and variance of the beta distribution and as a function of gamma-distributed random variables:

$$X \sim \Gamma(-\frac{E_Z(E_Z^2 - E_Z + V_Z)}{V_Z}, 1)$$

$$Y \sim \Gamma(\frac{(E_Z - 1)(V_Z + E_Z^2 - E_Z)}{V_Z}, 1)$$

$$Z = \frac{X}{X + Y} \sim Beta(\mu = E_Z, \sigma^2 = V_Z)$$

These formulas were used to implement the beta distribution in Netlogo and generalise over the distributions of continuous heterogeneous characteristics in the general model (H_i, α_i, w_i).

8.2 Generalising over the Joint Distribution of Agent Attributes

8.2.1 Constructing variables with pre-specified marginal distributions and pre-specified correlations

The following provides the formal reasoning underlying the approach to creating correlation between continuous variables in the simulation.

Let us say that we have two random-variables (i.e. random over the agent population): say, H and W. These have marginal probability density functions $p_H()$ and $p_W()$, which are specified directly. They also have associated cumulative distribution functions $F_H()$ and $F_W()$.

We also have a random variable H_{noise} which has marginal distribution $p_{H_{noise}}()$, and cumulative distribution function $F_{H_{noise}}()$ and whose Pearson correlation with H has been specified directly by defining H_{noise} as $H + \epsilon_H$ where ϵ_H is a random Gaussian noise variable with a standard deviation specified to create the desired correlation between H_{noise} and H (see 8.2.2).

We want to specify W such that it maintains the marginal distribution $p_W()$ but the joint distribution $p(H, W)$ has a non-linear correlation equal to the Pearson correlation of H and H_{noise}. To achieve this, we define W so that it retains its marginal distribution but has a perfect non-linear correlation with H_{noise}. We define W as follows:

$$W = F_W^{-1}([F_{H_{noise}}(H + \epsilon_H)])$$

Where F^{-1} is the inverse CDF of W.

8.2.2 Arbitrary correlations using random noise with a pre-specified variance

Theorem (Arbitrary correlations using random noise with a pre-specified variance)

Given:

$$y = x + \epsilon$$

Where x and ϵ are random independent variables and ϵ has an expected value of 0. It will be the case:

$$\rho_{x,y} = \frac{\sigma_x}{\sqrt{\sigma_x^2 + \sigma_\epsilon^2}}$$

Where $\rho_{x,y}$ is the Pearson product-moment correlation between x and y, and σ_x is the standard deviation of x, etc.

Proof (Arbitrary correlations using random noise with a pre-specified variance)

To prove this, we rely on the following previously established theorems regarding the properties of the expected value operator ($E[\cdot]$), the definition of the variance of a random variable in terms of expected value, and the definition of the Pearson correlation.

The person correlation can be defined:

$$\rho_{x,y} = \frac{E[(x - E[x])(y - E[y])]}{\sigma_x \cdot \sigma_y}$$

Where $E[\cdot]$ is the expected value operator, which has the following established properties:

3. E is distributive with respect to addition: $E(x) + E(y) = E(x + y)$
4. If x and y are independent, E is distributive with respect to multiplication: $x \perp y \rightarrow E(xy) = E(x) \cdot E(y)$
5. The expected value operator applied to a non-random variable returns that variable, e.g., $E(2) = 2, E(E(x)) = E(x)$.
6. Constants can be factored out of the expected value operator, such that $E(x \cdot E(x)) = E(x)^2$

Finally, we rely on the following definition of the variance of a random variable:

$$\sigma_x^2 = V[x] = E[x^2] - E[x]^2$$

Proof of the theorem depends on simplification and substitution within the denominator and the numerator in the definition of the Pearson correlation, for the case in which $y = x + \epsilon$.

$$E[(x - E[x])(y - E[y])]$$

$$= E[xy - xE[y] - yE[x] + E[x]E[y]]$$

$$= E(xy) - E(xE[y]) - E(yE[x]) + E(E[x]E[y])$$

$$= E(xy) - E(x)E(y) - E(x)E(y) + E(x)E(y)$$

$$= E(xy) - E(x)E(y)$$

$$sub. y = x + \epsilon \Longrightarrow \ldots = E[x(x + \epsilon)] - E[x]E[x + \epsilon]$$

$$= E[x^2] + E(x \cdot \epsilon) - E[x](E[x] + E[\epsilon])$$

$$= E[x^2] + E(x \cdot \epsilon) - E[x]^2 - E[x]E[\epsilon]$$

$$x \perp \epsilon \Longrightarrow \ldots = E[x^2] + E[x]E[\epsilon] - E[x]^2 - E[x]E[\epsilon]$$

$$= E[x^2] - E[x]^2 = \sigma_x^2$$

So, in the case where $y = x + \epsilon$, $E(\epsilon) = 0$ and $x \perp \epsilon$ (independence), the covariance of x and y will be equal to the variance of x.

Therefore, we can re-state the Pearson correlation between x and y as:

$$\rho_{x,y} = \frac{\sigma_x \cdot \sigma_x}{\sigma_x \cdot \sigma_y}$$

$$\Longrightarrow \rho_{x,y} = \frac{\sigma_x}{\sigma_y}$$

We can assume that σ_y is unknown, whereas σ_x and σ_ϵ are known. As such, it is helpful to re-write this as:

$$\Longrightarrow \rho_{x,y} = \frac{\sigma_x}{\sigma_y} = \frac{\sigma_x}{\sqrt{\sigma_x^2 + \sigma_\epsilon^2}}$$

This follows from the distributivity of the variance operator $V(\cdot)$ with respect to addition, under the independence of random variables:

$$x \perp \epsilon \Longrightarrow V(x+\epsilon) = V(x) + V(\epsilon)$$

This can be proved by substituting the definition of $y = x + \epsilon$ into the definition of the variance of y in terms of $E[\cdot]$, and then simplifying. The distributivity of the variance operator with respect to addition depends on the independence of ϵ and x, but doesn't require that $E(\epsilon)$ or $E(x)$ be equal to 0 (as other parts of the proof do).

As such, we can note that:

$$\sigma_y = \sqrt{V(y)} \Longrightarrow \sigma_y = \sqrt{V(x) + V(\epsilon)} = \sqrt{\sigma_x^2 + \sigma_\epsilon^2}$$

Application (Arbitrary correlations using random noise with a pre-specified variance)

The intended application of the theorem is the generation of random variables which have an arbitrary degree of correlation with some prior variable. Rearranging the theorem (assuming all values are positive and greater than 0) above shows that this can be achieved using the following formulation:

$$y = x + \epsilon$$

Where:

$$\sigma_\epsilon = \frac{\sqrt{\sigma_x^2 - \rho_{x,y} \cdot \sigma_x^2}}{\rho_{x,y}}$$

With the desired $\rho_{x,y}$ chosen arbitrarily. ϵ can have any probability density as long as $E(\epsilon) = 0$.

8.3 Parameterizing 'Social Reach' in the Hamill and Gilbert Social Circle Algorithm in terms of the Expected Connectivity of Actors

To build the social reference and decision-maker reference networks, a simple version of Hamill and Gilbert's (2009) spatial network algorithm was used. This algorithm implements the following steps:

1. Construct a population of N agents
2. Distribute the agents randomly across a 2D space which wraps at its edges
3. Connect all agents who are within a radius of r of one-another

The raw parametrization of the algorithm is in terms of r; however, this is largely an unintuitive parameter, and moreover, it is insensitive to the concentration of actors in the space - which will affect overall network density. Instead, we would like to re-parameterise the algorithm in terms of the expected density of the resulting network. In order to do this, we will need to express r in terms of the expected number of connections of each agent in the network.

First, note that actor i connects to actor j *iff* the location of actor j is within r of actor i. Since all positions for j within the 2D space are equally likely, the probability that i is within r of j is the proportion of points in space for which this will occur. This is simply the ratio of the area of the radius r and the area of the wrapped 2D space, which is:

$$P_{ij} = \frac{\pi r^2}{A_t}$$

Where A_t is the area of the 2D space and P_{ij} is the probability that actor i will connect to actor j, for $j \neq i$.

Since this probability is independent for all other actors in the agent population, we can expect the number of actors connected to i to be Binomial distributed, with $N - 1$ independent trials, each with probability $\frac{\pi r^2}{A_t}$ of success (connection). Here we are simply interested in the expected value for i, which will be $(N-1) \cdot \frac{\pi r^2}{A_t}$. We will call this figure (the expected connectivity of each actor) E_c.

We can then parameterise the social circle algorithm by defining r in terms of E_c:

$$E_c = (N-1) \cdot \frac{\pi r^2}{A_t} \Longrightarrow r = \sqrt{\frac{A_t \cdot E_c}{\pi \cdot (N-1)}}$$

This formula is used in the model, with E_c controlled by parameter $\mu_{Rsocial}$ (in the case of the social reference network), to determine the social-reach of actors, and construct a network with the desired connectivity.

8.4 Derivation of the Probability that 'Decision-Maker' Agents will Update their Cutting decisions in a Single Time-Step

The calculation of the probability that, in a given time-step, a 'decision-maker' agent would need to make a 'final decision' about the cutting status of one of the girls in their household proceeded as follows. Relevant empirical figures for b and c were taken from an international survey of household characteristics (United Nations Department of Economics and Social Affairs, 2017: 16), assuming that, on average, 50% of children under 15 are girls.

Let a = *the number of households in the community, where* n *is the number of adults*

$$a = \frac{n}{6}$$

Let b = *the proportion of households with children*

$$b = 0.84$$

Let c = *the average number of female children (aged under 15) per household*

$$c = 2.25$$

Let d = *the average number of children in the community*

$$d = a \cdot b \cdot c = \frac{n}{6} \cdot 0.84 \cdot 2.25$$

Let l = *the number of decisions to be taken each year such that, on average, all girls are cut by age 15 if they are to be cut* [4]

$$l = \frac{d}{15}$$

Let m = *the number of final decision-makers in the community, assuming one per-household*

4 Kandala and Shell-Duncan (2019: Table 1) report that less than 1% of girls are cut *after* the age of 15.

$$m = a = \frac{n}{6}$$

Assuming x *is the number of 'time-points'* (t_k) *related to FGM each year, let* $Pr(decision|t_k = k)_i$ *be the probability that decision-maker* i *takes a decision about one-of their girls in a single time-point*

$$P(decision|t_k = k)_i = \frac{l}{x \cdot m}$$